

NAVAL POSTGRADUATE SCHOOL

Monterey, California



THESIS

INTERNETWORKING: NPS ATM LAN

by

Dale Michael Courtney

September, 1996

Thesis Advisor:
Associate Advisor:

Don Brutzman
Rex Buddenberg

Approved for public release; distribution is unlimited.

19970107 134

DTIC QUALITY INSPECTED 1

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.				
1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE September 1996	3. REPORT TYPE AND DATES COVERED Master's Thesis		
4. TITLE AND SUBTITLE: INTERNETWORKING: NPS ATM LAN		5. FUNDING NUMBERS		
6. AUTHOR Courtney, Dale Michael				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey CA 93943-5000		8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSORING/MONITORING AGENCY REPORT NUMBER		
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.				
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT The objective of this research is to create, build, and test an electronic information infrastructure at NPS based on ATM cell relay, and to lay the groundwork for future ATM work at NPS. One aspect of this research is to critique ATM as a future networking technology for DoD and the U.S. Navy. This research demonstrates five fatal flaws of ATM with respect to the military environment. First, there is the interoperability between switches. There is no way to guarantee communication between switches. Second, there is ATM's incompatibility with IP. There is no native way to multicast with ATM. Overcoming the multicasting problem is probably the greatest ATM problem to solve, and on-going research has yet to find a native ATM solution to this problem. Third, there is ATM's inflexibility to change. Myriad long-haul problems exist. Forth, there is the human factor. The "expertise" that exists in the ATM field is nominal due to the immaturity of the technology. Fifth, there is the crossover problem. The crossover system from primary to backup mechanisms must be reliable. ATM has not solved the problem of crossover. If a connection is broken, there is no standby connection waiting to immediately take over; and this scenario is exacerbated in the already problematic multicast situation. Before DoD becomes too committed to ATM, these five issues need to be explicitly and fully resolved.				
14. SUBJECT TERMS Computer Network Architecture; Asynchronous Transfer Mode (ATM); Local-Area Network (LAN); Network Topology; Multicast.			15. NUMBER OF PAGES 211	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL	

Approved for public release; distribution is unlimited.

INTERNETWORKING: NPS ATM LAN

Dale Michael Courtney
Lieutenant Commander, United States Navy
B.S., University of Colorado, 1982

Submitted in partial fulfillment
of the requirements for the degree of

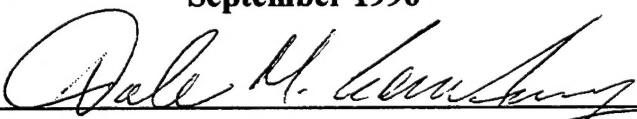
**MASTER OF SCIENCE IN
INFORMATION TECHNOLOGY MANAGEMENT**

from the

NAVAL POSTGRADUATE SCHOOL

September 1996

Author:



Dale M. Courtney

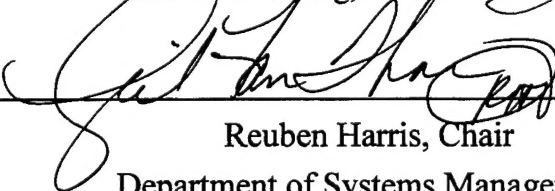
Approved by:



Don Brutzman, Thesis Advisor



Rex Buddenberg, Associate Advisor



Reuben Harris, Chair

Department of Systems Management

ABSTRACT

The objective of this research is to create, build, and test an electronic information infrastructure at NPS based on ATM cell relay, and to lay the groundwork for future ATM work at NPS.

One aspect of this research is to critique ATM as a future networking technology for DoD and the U.S. Navy. This research demonstrates five fatal flaws of ATM with respect to the military environment. First, there is the interoperability between switches. There is no way to guarantee communication between switches. Second, there is ATM's incompatibility with IP. There is no native way to multicast with ATM. Overcoming the multicasting problem is probably the greatest ATM problem to solve, and on-going research has yet to find a native ATM solution to this problem. Third, there is ATM's inflexibility to change. Myriad long-haul problems exist. Forth, there is the human factor. The "expertise" that exists in the ATM field is nominal due to the immaturity of the technology. Fifth, there is the crossover problem. The crossover system from primary to backup mechanisms must be reliable. ATM has not solved the problem of crossover. If a connection is broken, there is no standby connection waiting to immediately take over; and this scenario is exacerbated in the already problematic multicast situation. Before DoD becomes too committed to ATM, these five issues need to be explicitly and fully resolved.

TABLE OF CONTENTS

I.	INTRODUCTION	1
A.	PROBLEM DEFINITION	1
B.	MOTIVATION	1
1.	ATM at NPS	2
2.	Regional and National Impetus	2
3.	Uses for the Navy	2
4.	Uses for DoD	3
C.	THESIS ORGANIZATION	3
II.	RELATED WORK	5
A.	INTRODUCTION	5
B.	IIRG	5
1.	Network Monitoring and Performance Evaluation for Global Networks	6
2.	Computer Security	7
3.	Multicast and ATM Network Prerequisites for Distance Learning	8
4.	MBone Monitoring	9
5.	Economical Storage and Retrieval of Digital Audio and Video for Distance Learning	9
6.	Wide-area Network (WAN) for K-12 Schools	10
C.	REGIONAL AND NATIONAL INTERESTS	11
1.	BayNet and BayLink	11
2.	CalREN	11
3.	BAGNet	12
4.	I-WAY	13
5.	IEEE Computer Graphics and Applications	14
D.	U.S. MILITARY INTERESTS	15
1.	U.S. Navy	15
2.	DoD	16
E.	SUMMARY	17
III.	PROBLEM STATEMENT	19
IV.	HARDWARE METHODOLOGY	21
A.	INTRODUCTION	21
B.	PRESENT CAMPUS FIBER TOPOLOGY	21
C.	DESIGNING THE FUTURE ATM NETWORK	22
D.	CHRONOLOGY OF EVENTS	25

1.	Stand Alone	26
2.	Hooking up Two Workstations Peer-to-Peer	28
3.	Hooking up a Single ATM Switch to Two Workstations	29
4.	Hooking up Two ATM Switches	30
5.	Connecting to BayNet	31
E.	FUTURE DESIRED CONFIGURATION	37
F.	SUMMARY	37
V.	SOFTWARE METHODOLOGY	41
A.	INTRODUCTION	41
B.	SIMPLE ATM SUMMARY	41
C.	ATM PROTOCOL REFERENCE MODEL	42
1.	The Physical Layer	44
2.	The ATM Layer	50
a.	ATM Cell Format	50
b.	Interfaces	53
(1)	NNI.	54
(2)	UNI.	57
c.	Cell Routing	59
3.	ATM Adaptation Layer (AAL)	61
D.	ATM SIGNALING	67
1.	Connection Management	68
2.	Connection Types	75
3.	Uni 3.0	78
4.	NSAP	81
E.	SUMMARY	83
VI.	EXPERIMENTAL RESULTS	85
A.	INTRODUCTION	85
B.	SOFTWARE TEST DESCRIPTION	85
C.	CONFIGURATIONS	86
1.	Stand Alone	86
2.	Hooking up Two Workstations Peer-to-peer	87
3.	Hooking up a Single ATM Switch to Two Workstations	88
4.	Hooking up Two ATM Switches	89
5.	Connecting to BayNet	90
D.	FAILURE TO SUSTAIN BAYNET CONNECTIVITY	93
E.	CONCLUSION	93
VII.	CONCLUSIONS AND RECOMMENDATIONS	95
A.	INTRODUCTION	95
B.	CONCLUSIONS	95
C.	FUTURE PLANS	96

1.	Proposed NPS ATM LAN Extensions	96
2.	Changes in the CalREN grant	97
D.	BARRIERS, FUTURE WORK, AND RECOMMENDATIONS	98
1.	Barriers	98
2.	Future Work	99
a.	Multicasting	99
b.	Reliability	100
c.	ATM over Satellite	101
3.	Recommendations for Future Work	101
E.	RECOMMENDATIONS TO DISA AND THE U.S. NAVY	102
1.	Interoperability between switches	103
2.	Incompatibility with IP	103
3.	Inflexibility to change	103
4.	Individuals	104
5.	Crossover	104
F.	SUMMARY	104
	LIST OF REFERENCES	109
	APPENDIX A. ABBREVIATIONS AND DEFINITIONS	117
	APPENDIX B. FORE CARD SPECIFICATIONS	139
A.	INTRODUCTION	139
B.	HARDWARE OVERVIEW	139
C.	SOFTWARE OVERVIEW	139
D.	VMA-200 ARCHITECTURE	141
E.	HARDWARE REQUIREMENTS	141
F.	SOFTWARE REQUIREMENTS	141
G.	FIBER OPTIC CABLE SPECIFICATIONS	141
H.	TECHNICAL SPECIFICATIONS	141
	APPENDIX C. CISCO A-100 SWITCH SPECIFICATIONS	145
A.	INTRODUCTION	145
B.	INTERFACE TYPES	145
C.	FUNCTIONAL OVERVIEW	150
1.	Cell Switching Function	150
2.	Traffic Control Functions	151
a.	Policing Control	151
b.	Congestion Control	152
c.	Priority Control	152
3.	Connection Control	153
D.	EXPANDABLE ATM OUTPUT BUFFER MODULAR SWITCH (ATOMSW)	155

1.	MUX/DMUX	156
2.	ATOMSW	158
a.	ATOMSW Cell Switching Function	159
b.	ATOMSW Congestion Control	160
c.	ATOMSW Priority Control	161
APPENDIX D. FORE CARD CONFIGURATION		163
A.	INTRODUCTION	163
B.	STAND ALONE	163
C.	HOOKING UP TWO WORKSTATIONS PEER-TO-PEER	164
D.	HOOKING UP A SINGLE ATM SWITCH TO TWO WORKSTATIONS	166
E.	HOOKING UP TWO ATM SWITCHES	167
F.	CONNECTING TO BAYNET	167
APPENDIX E. CONFIGURING THE CISCO A-100 SWITCH		169
A.	INTRODUCTION	169
B.	STAND ALONE	169
C.	HOOKING UP TWO WORKSTATIONS PEER-TO-PEER	169
D.	HOOKING UP A SINGLE ATM SWITCH TO TWO WORKSTATIONS	170
1.	Configuring a PVC	172
2.	Configuring an SVC	173
E.	HOOKING UP TWO ATM SWITCHES	174
F.	CONNECTING TO BAYNET	178
APPENDIX F. RETRIEVING ONLINE THESIS		181
INITIAL DISTRIBUTION LIST		183

LIST OF FIGURES

Figure 1. I-Way Prototyping Testbed, after [I-WAY, 1995].	14
Figure 2. Applications of NRL's ATM Network, after [ATDNet, 1994].	15
Figure 3. NPS Campus Backbone Topology, after [Romo, 1995].	22
Figure 4. ATM resources, after [Advisory Group, 1995].	23
Figure 5. Activation Priorities, after [Advisory Group, 1995].	24
Figure 6. Requirements to Meet the Goal, after [Advisory Group, 1995].	24
Figure 7. Network Administrator Roy Romo's Concerns, after [Romo, 1995].	25
Figure 8. Putting the Concerns to Rest.	26
Figure 9. Stand Alone Configuration.	27
Figure 10. Peer-to-Peer Configuration.	28
Figure 11. Multi-homing of ATM and Ethernet.	29
Figure 12. Single-Switched ATM LAN.	30
Figure 13. Dual-Switched ATM LAN.	30
Figure 14. NPS ATM LAN Connection to BayNet Cloud.	31
Figure 15. BayNet Configuration, after [Arul, 1996].	32
Figure 16. UCSC's ATM Configuration, after [Arul, 1996].	33
Figure 17. Sprint's and PacBell's Switches, after [Arul, 1996].	34
Figure 18. InterLATA provisions of the 96 Act , after [Sapronov, 1996].	36
Figure 19. NPS ATM LAN — IP and ATM connections.	37
Figure 20. Final Desired Configuration, after [Advisory Group, 1995].	38
Figure 21. Four Major Assumptions Behind ATM, after [Partridge, 1994].	41
Figure 22. Three Domains (Planes) of ATM Activity, after [Feibel, 1995].	43
Figure 23. Three layers created in the ATM Hierarchy, after [Feibel, 1995].	43
Figure 24. Relationship between SONET and SDH, from [Symborski, 1995].	45
Figure 25. ATM Physical Layer, after [Varma, 1995].	48
Figure 26. Physical Interfaces Allowed by the ATM Forum, after [Feibel, 1995].	49
Figure 27. ATM Cell Format, from [Partridge, 1994].	51
Figure 28. Cell-size Tradeoffs, after [Varma, 1995].	52
Figure 29. UNI's Two-fold Purpose, after [Partridge, 1994].	53
Figure 30. ATM Header at NNI, from [Partridge, 1994].	54
Figure 31. Payload-type Field, from [Varma, 1995].	55
Figure 32. ATM Header at UNI, from [Partridge, 1994].	57
Figure 33. Channel Identifiers, after [Alles, 1995].	59
Figure 34. Traffic or Service Proposed by CCITT, after [Varma, 1995].	62
Figure 35. Four B-ISDN Service Classes, after [Siu, 1994].	63
Figure 36. Goals of AAL5 (SEAL), after [Partridge, 1994].	64
Figure 37. AAL Sublayers, after [Feibel, 1995].	64
Figure 38. ATM Adaptation Layer (AAL), after [Varma, 1995].	65
Figure 39. AAL 5 SAR Format, from [Partridge, 1994].	66
Figure 40. AAL5 Convergence Sublayer Format, from [Partridge, 1994].	66

Figure 41. Two Fundamental Types of ATM Connections, after [Alles, 1995].	67
Figure 42. Q.93B Signaling Connection Setup, from [Partridge, 1994].	70
Figure 43. Sample ATM Switching Network, after [Ebrahim, 1992].	71
Figure 44. ATM Connection Types, after [Alles, 1995].	75
Figure 45. Proposed ATM Multicasting Methods, after [Alles, 1995].	77
Figure 46. Three Formats of Private ATM Addressing, after [Alles, 1995].	82
Figure 47. ICD ATM Format, from [Alles, 1995].	83
Figure 48. Stand Alone.	86
Figure 49. Peer-to-Peer.	88
Figure 50. Single-Switched.	89
Figure 51. Dual-Switched.	90
Figure 52. BayNet Configuration.	91
Figure 53. X-Window Applications.	92
Figure 54. NPS ATM LAN Connection to BayNet Cloud.	96
Figure 55. Final Desired Configuration, after [Advisory Group, 1995].	97
Figure 56. Criteria for Survivability and Fault Tolerance, from [Buddenberg, 1995].	100
Figure 57. ATM "Show-stoppers."	103
Figure 58. VMA-200 VMEbus Adapter, from [Fore, 1995].	140
Figure 59. VMA-200 Cell Processing Architecture, from [Fore, 1995].	142
Figure 60. Systems That Can Use the VMA-200, from [Fore, 1995].	142
Figure 61. Cisco A-100 Switch Features, from [Cisco, 1995].	148
Figure 62. Cisco A-100 Switch Primary Functions, from [Cisco, 1995].	150
Figure 63. Traffic Measurement Period, from [Cisco, 1995].	152
Figure 64. Cell Type Priority, after [Cisco, 1995].	153
Figure 65. A-100 VPI/VCI Address Space, from [Cisco, 1995].	154
Figure 66. A-100 PVCs and SVCs Possible, from [Cisco, 1995].	154
Figure 67. PVC Elements Controlled by the A-100, from [Cisco, 1995].	155
Figure 68. ATOMSW Block Diagram, from [Cisco, 1995].	156
Figure 69. MUX/DMUX Card Block Diagram, from [Cisco, 1995].	157
Figure 70. ATOMSW Block Diagram, from [Cisco, 1995].	159
Figure 71. ATOMSW Congestion Control Threshold Values, from [Cisco, 1995]. . .	160
Figure 72. ATOMSW Threshold Values, from [Cisco, 1995].	161
Figure 73. ATOMSW Output-Buffer Threshold Values, from [Cisco, 1995].	161
Figure 74. Stand Alone Configuration.	164
Figure 75. Peer-to-Peer Configuration.	165
Figure 76. Single-Switched ATM LAN.	166
Figure 77. Dual-Switched ATM LAN.	167
Figure 78. NPS ATM LAN Connection to BayNet Cloud.	168

Figure 79. Stand Alone Configuration.	170
Figure 80. Peer-to-Peer Configuration.	171
Figure 81. Single-Switched ATM LAN.	171
Figure 82. Dual-Switched ATM LAN.	175
Figure 83. Output of the initial show atmsig command.	177
Figure 84. Output of the show atmsig command after resetting User/Network (U/N) status for line 0.	178
Figure 85. NPS ATM LAN Connection to BayNet Cloud.	179

LIST OF TABLES

Table 1. Asynchronous Data Rates, after [Freeman, 1991; Partridge, 1994].	46
Table 2. SONET/SDH Transmission Rates, after [Freeman, 1991; Partridge, 1994]. .	47
Table 3. Cisco A-100 Interface Types, from [Cisco, 1995].	49
Table 4. Payload Type Values, after [Partridge, 1994; Varma, 1995].	55
Table 5. AAL Types, from [Varma, 1995].	62
Table 6. Switch B's Routing Table. Boldface values show table entries for example shown in Table 5.	73
Table 7. Example VC Table Mappings, after [Ebrahim, 1992].	74
Table 8. OSI Layer Descriptions for Packet.	130
Table 9. Fiber Optic Cable Specifications for the VMA-200, from [Fore, 1995]. . .	143
Table 10. Technical Specifications of the Fore VMA-200, after [Fore, 1995].	144
Table 11. Cisco A-100 Switch Interface Types, from [Cisco, 1995].	146
Table 12. ITU-T Specifications Supported by the Cisco A-100, from [Cisco, 1995].	147
Table 13. Cisco A-100 Switch Specifications, from [Cisco, 1995].	149
Table 14. Conditions in which FIFO Queues Are Served, from [Cisco, 1995].	158

LIST OF ABBREVIATIONS, ACRONYMS, AND SYMBOLS

The following list of abbreviations is a version of Fore's acronym handbook. The abridgement is based on terms used in this thesis. The entire glossary is available at <http://www.fore.com/html/atm-edu/acronyms.html>

AAL:	ATM Adaptation Layer
ABR:	Available Bit Rate
ADPCM:	Adaptive Differential Pulse Code Modulation
AII:	Active Input Interface (Used in UNI PMD specs for Copper/Fiber)
ANSI:	American National Standards Institute
AOI:	Active Output Interface (Used in UNI PMD specs for Copper/Fiber)
API:	Applications Programming Interface
ARP:	Address Resolution Protocol
ASN.1:	Abstract Syntax Notation One
ATE:	ATM Terminating Equipment (SONET)
ATM:	Asynchronous Transfer Mode
BER:	Basic Encoding Rules (ASN.1)
BER:	Bit Error Rate (link quality specification/testing)
B-ISDN:	Broadband Integrated Services Digital Network
BOM:	Beginning of Message
BRI:	Basic Rate Interface
CAC:	Connection Admission Control
CBR:	Constant Bit Rate
CCITT:	Consultative Committee on International Telephone and Telegraph (now ITU)
CDV:	Cell Delay Variation
CIR:	Committed Information Rate
CLP:	Cell Loss Priority
CLR:	Cell Loss Ratio
CPCS:	Common Part Convergence Sublayer
CPI:	Common Part Indicator
CRC:	Cyclic Redundancy Check
DE:	Discard Eligibility
DLCI:	Data Link Connection Identifier
DMA:	Direct Memory Access

DNS:	Domain Name Service
DQDB:	Distributed Queue Dual Bus
DS-1:	Digital Signal Level 1
DS-3:	Digital Signal Level 3
DSU:	Data Service Unit
DXI:	Digital Exchange Interface
DXI:	Data Exchange Interface
E-1:	European Signal Level 1
E-3:	European Signal Level 3
EISA:	Extended Industry Standard Architecture
ELAN:	Emulated LAN (ATM Forum LANE)
EOM:	End of Message
FCS:	Frame Check Sequence
FCVC:	Flow Controlled Virtual Circuit
FDDI:	Fiber Distributed Data Interface
FTP:	File Transfer Protocol
GFC:	Generic Flow Control
HEC:	Header Error Control
HPPI:	High Performance Parallel Interface
HSSI:	High Speed Serial Interface
ICI:	Inter-Carrier Interface
IE:	Information Element
IEEE:	Institute of Electrical and Electronic Engineers
IETF:	Internet Engineering Task Force
ILMI:	Interim Local Management Interface
IP:	Internet Protocol
IPX:	Internet Packet Exchange
ISDN:	Integrated Services Digital Network
ISO:	International Standards Organization
ITU:	International Telecommunications Union (formerly CCITT)
JPEG:	Joint Photographic Experts Group
LAN:	Local Area Network
LANE:	LAN Emulation (ATM Forum)
LATA:	Local Access Transport Area
LE:	Link Encapsulation
LEC:	LAN Emulation Client or Local Exchange Carrier
LENNI:	LAN Emulation NNI
LES:	LAN Emulation Service
LIS:	Logical IP Subnet
LLC:	Logical Link Control
LSB:	Least Significant Bit

LUNI:	LAN Emulation UNI
MAC:	Media Access Control
MAN:	Metropolitan Area Network
MBS:	Maximum Burst Size
MCR:	Minimum Cell Rate
MID:	Message Identifier
MIB:	Management Information Base
MPEG:	Motion Picture Experts Group
MPOA:	MultiProtocol Over ATM
MSB:	Most Significant Bit
MTU:	Maximum Transfer Unit
NBMA:	Non-Broadcast Multiple Access
NDIS:	Network Driver Interface Specification
NHRP:	Next Hop Resolution Protocol
NIU:	Network Interface Unit
NLPID:	Network Layer Protocol ID
NNI:	Network-to-Network Interface
NSAP:	Network Layer Service Access Point
NTSC:	National Television Standards Committee
OAM:	Operations, Administration, and Maintenance
OCD:	Out-of-Cell Delineation (UNI 3.0)
OC-n:	Optical Carrier-n
ODI:	Open Data Link Interface
OS:	Operating System
OSI:	Open Systems Interconnect
PCR:	Peak Cell Rate
PDU:	Packet Data Unit
PING:	Packet Internet Groper
PLCP:	Physical Layer Convergence Protocol
PMD:	Physical Medium Dependent
PMP:	Point to Multipoint
P-NNI:	Private Network-to-Network Interface or Private Network Node Interface
POH:	Path Overhead (SONET)
POI:	Path Overhead Indicator
PPP:	Point-to-Point Protocol
PRI:	Primary Rate Interface
PSTN:	Public Switched Telephone Network
PT:	Payload Type
PTE:	Path Terminating Equipment (SONET)
PTI:	Payload Type Identifier
PVC:	Permanent Virtual Circuit

QoS:	Quality of Service
RAM:	Random Access Memory
RBOC:	Regional Bell Operating Company
RFC:	Request For Comment
RIP:	Routing Information Protocol
ROLC:	Routing Over Large Clouds
ROM:	Read-Only Memory
SAA:	Systems Applications Architecture
SAAL:	Signaling ATM Adaptation Layer
SAP:	Service Access Point
SAR:	Segmentation and Reassembly
SCSI:	Small Computer Systems Interface
SCR:	Sustainable Cell Rate
SDH:	Synchronous Digital Hierarchy
SDLC:	Synchronous Data Link Control
SDU:	Service Data Unit
SEAL:	Simple and Efficient Adaptation Layer
SFD:	Start of Frame Delimiter
SIMM:	Single In-line Memory Module
SIR:	Sustained Information Rate
SMDS:	Switched Multimegabit Data Service
SNA:	Systems Network Architecture
SNAP:	Sub-Network (Area Protocol) Attachment Point
SNMP:	Simple Network Management Protocol
SONET:	Synchronous Optical Network
SPANS:	Simple Protocol for ATM Network Signaling
SRTS:	Synchronous Residual Time Stamp
SS:	Switching System
SS7:	Signaling System #7
SSCS:	Service Specific Convergence Sublayer
SSCOP:	Service Specific Connection Oriented Protocol
STDM:	Statistical Time Division Multiplexing
STE:	Section Terminating Equipment (SONET)
STM:	Synchronous Transport Module
STP:	Shielded Twisted Pair
STS-n:	Synchronous Transport Signal-n
SVC:	Switched Virtual Circuit
TA:	Terminal Adapter
TAXI:	Transparent Asynchronous Transmitter/Receiver Interface
TC:	Transmission Convergence
TCP:	Transmission Control Protocol

TCP/IP:	Transmission Control Protocol/Internet Protocol
TDM:	Time Division Multiplexing
TDS:	Time Division Switching
TS:	Telecommunication Standardization sector
UBR:	Unspecified Bit Rate
UDP:	User Datagram Protocol
UNI:	User-to-Network Interface
UPC:	Usage Parameter Control
UTP:	Unshielded Twisted Pair
VBR:	Variable Bit Rate
VC:	Virtual Circuit
VCC:	Virtual Channel Connection
VCFC:	Virtual Channel Flow Control
VCI:	Virtual Channel Identifier
VCL:	Virtual Channel Link (UNI 3.0)
VP:	Virtual Path
VPC:	Virtual Path Connection
VPCI:	Virtual Path Connection Identifier
VPI:	Virtual Path Identifier
VPL:	Virtual Path Link
VPN:	Virtual Private Network
VPT:	Virtual Path Terminator (UNI 3.0)
WAN:	Wide Area Network

ACKNOWLEDGMENTS

The author acknowledges the support, encouragement, and prayers of his wife, Carla, and his children Chantelle, Caleb, Joshua, and Joel. I could not have made it through school without their constant assistance, perseverance, and tenacity along side of me.

The author acknowledges the assistance of his good friends in the Information Infrastructure Research Group — Jim Cummiskey, Ron Dennis, Evan Edwards, Ridvan Erdogan, Steve Graves, Lauren Mihlon, Murat Tamer, and Mike Tiddy. Without their involvement in this research, I would never have finished.

The author sincerely appreciates the encouragement of his thesis advisors: Don Brutzman, who taught me that you really **can** do anything, even if you don't understand it all; and Rex Buddenberg, whose depth and breadth of knowledge in networking I hope to achieve in my own life.

And to my best friends: Ron, Steve, and Kevin — I wouldn't have graduated without you all!

I. INTRODUCTION

A. PROBLEM DEFINITION

The purpose of this research is to develop an asynchronous transfer mode (ATM) local-area network (LAN) at the Naval Postgraduate School (NPS). The purpose of this LAN is to perform research into the capabilities of ATM, to effectively integrate ATM with standard internetworked LAN technologies such as FDDI and Ethernet, with a regional ATM wide-area network (WAN), and to enable various research initiatives to demonstrate effective unicast and multicast connectivity at NPS using ATM technology.

B. MOTIVATION

The motivation behind building the NPS ATM LAN is fourfold. First, the ATM LAN allows NPS students to research the capabilities of ATM, especially with respect to internetworking. Secondly, there is regional and national impetus to test the huge bandwidth and low-latency capabilities of ATM. Thirdly, the U.S. Navy is heavily involved in ATM research, development, and operations at the Naval Research Laboratory (NRL). Fourthly, the DoD is heavily involved in ATM initiatives and experimentation at the Defense Information Systems Agency (DISA).

Student research at NPS will add to the Navy's and DoD's understanding of ATM advantages and disadvantages. NPS is examining the benefits and problems with ATM technology in order to provide the Navy and DoD with the feedback it needs to correctly evaluate the risk of transferring to a cutting edge technology such as ATM.

1. ATM at NPS

The ATM LAN allows NPS to research the capabilities of ATM: high bandwidth, low latency, network monitoring, network security, multicast and Multicast Backbone (MBone) monitoring, and on-line storage and retrieval. This related research is detailed in Chapter II.

2. Regional and National Impetus

Pacific Bell (PacBell) established the California Research and Education Network (CalREN) to stimulate the development of new applications for high-speed data communications services. The CalREN program is funded at a level of \$25 million statewide, so there are research opportunities available regionally. These are also discussed in detail in Chapter II and include BayNet, BayLink, CalREN, BAGNet, and national interests such as the I-WAY and Virtual Chesapeake Bay.

3. Uses for the Navy

The Navy operates an ATM network at the Naval Research Laboratory (NRL) which serves as a test bed and supports mission traffic. The New Attack Submarine (NSSN) architecture subsystem will utilize ATM for providing connectivity services to the subsystems comprising the NSSN C3I system [NSSN, 1995].

NPS is examining the benefits and problems with ATM technology in order to provide NRL and other Navy labs with the feedback needed to correctly evaluate the strengths and weaknesses of ATM when used in shipboard or shore environments.

4. Uses for DoD

The Defense Information Systems Agency (DISA) has focused on ATM as being the underlying network of the future. NPS is examining the benefits and problems with ATM technology in order to provide DISA and other Department of Defense (DoD) agencies with the feedback needed to correctly evaluate the risk of transferring to a cutting-edge technology such as ATM.

C. THESIS ORGANIZATION

This thesis is organized in the following manner: Chapter II describes ATM-related work at NPS regionally, globally, and related to the Navy and DoD. Chapter III details the actual problem researched in this thesis. Chapters IV and V discuss the hardware and software methodologies of building the ATM LAN, respectfully. Chapter VI provides the experimental results. Chapter VII furnishes conclusions and recommendations for future work.

There are six appendices to this thesis that provide a list of abbreviations and definitions, detail the specifications of the network interface cards (NICs) and of the ATM switches used, provide the details of configuring the NICs and switches, and provide directions how to retrieve this thesis online via the Internet.

II. RELATED WORK

A. INTRODUCTION

This chapter discusses work related to the NPS ATM LAN. It begins with local research occurring at NPS and then branches out to regional, national, and defense interests.

B. IIRG

The Information Infrastructure Research Group (IIRG) is an action team of thesis research students at the Naval Postgraduate School. The shared objective is to examine how to connect everyone to everything, focusing on the globally shared resource of the Internet. IIRG involvement is broad: internetworking K-12 schools and educational institutions, research into large-scale virtual environments (LSVEs), extending high-bandwidth low-latency connectivity using Asynchronous Transfer Mode (ATM), providing global many-to-many real-time audio/video connectivity using the Multicast Backbone (MBone), and defining Internet Protocol (IP) over seawater (IP/SW). Its narrower goals are to internetwork the U.S. Navy using a digital library, Internet to sea (SeaNet), and other projects.

The following is a list of recently completed NPS theses that relate to this work on building the NPS ATM LAN.

1. Network Monitoring and Performance Evaluation for Global Networks

Internetworking: Network Monitoring Aspects of High Bandwidth, Low Latency

Global Networks [Edwards, 1996] is a thesis that addresses the common lack of capability to monitor networks, particularly large networks and internetworks. To monitor a network means to have the capacity to determine the route taken by a communication, to display the time that it required, and to determine what percentage (if any) of the communication was lost. The complexity of network monitoring increases with each link in the connection. Proper development and evaluation of IIRG theses outlined in this section is dependent on network monitoring capabilities.

Monitoring capabilities do exist but with significant hindrances. Commercial tools are prohibitively expensive, while public domain software tools are text based, command-line driven, and cryptic. Neither commercial nor public domain tools represent a viable option for most network users. With proper automation and integration, public domain software tools can deliver accurate and timely information on network status and performance.

The internetworked NPS ATM LAN is the initial arena to test the monitoring tools. Following successful development of these tools in the local environment, they will be tested on the Monterey Bay Network (BayNet) using both Frame Relay and ATM backbones. Testing on a national level will follow BayNet evaluation. This will enhance collaboration with national research partners on the I-WAY [I-WAY, 1995], Old

Dominion University in Norfolk Virginia [Wheless, 1996], and Naval Undersea Warfare Center (NUWC) in Newport Rhode Island [NRL, 1996].

Online scripts are being constructed which monitor active BayNet ATM virtual circuits, construct home pages at regular intervals to show network status and notify system administration personnel of initial failures and when hosts again become available, archive system status automatically, and provide interactive means of network testing with simple home page interface. This effort was partially inspired by public domain source code for the network management project for the Bay Area Gigabit test bed Network (BAGNet). Current work includes reconfiguring BayNet ATM logical topology. Future work will extend this network monitoring model using other public domain software tools.

2. Computer Security

Internetworking: IP/ATM LAN Security [Dennis, 1996] is a thesis that addresses the security requirements of an IP/ATM LAN at the Naval Postgraduate School with an emphasis on internetworking and remote access security.

Originally the exclusive domain of scientists and researchers, the Internet is used today by millions of users around the world. Common tasks include transferring files, accessing data bases, conversing via one or more of the thousands of bulletin board systems, conducting interactive distance education, and conducting commercial business. As the number of Internet users has grown, so too has the number of computer security incidents.

Many believe that new connection-oriented network technologies such as ATM will reduce the threats to network security by eliminating unexpected connections. Regardless of the physical technology and protocols used for the Internet backbone, internetworking trends will continue. Correspondingly, computer and network security are also growing in importance. Personal data, credit card numbers, proprietary business information, and sensitive research data must be protected both while it is stored in host systems and while being transmitted over the network.

This thesis focuses on the vulnerabilities associated with internetworking a combined Internet Protocol (IP) Ethernet and ATM LAN at the Naval Postgraduate School with other networks over a wide area and discusses the Kerberos authentication and authorization protocol implemented to reduce the vulnerability of unauthorized access.

3. Multicast and ATM Network Prerequisites for Distance Learning

Internetworking: Multicast and ATM Network Prerequisites for Distance Learning [Tamer, 1996] is a thesis that examines how the Mbone and Mbone tools can be used most effectively. It investigates all requirements for using high-bandwidth, low-latency ATM infrastructure to enable economical near-real-time model interaction and scientific visualization of Virtual Environments (VEs) at geographically distributed sites.

The goal is to have good audio/video (A/V) quality through common Internet connections such as 128 Kbps Frame Relay, to have a simple, cost-effective configuration

for MBone A/V recording so that schools can afford MBone for their educational purposes, to facilitate a short period of personnel training, and to be able to use the MBone for real-time A/V over ATM.

4. MBone Monitoring

Internetworking: Implementation of Multicast and MBone Over Frame Relay Network [Erdogan, 1996] documents the implementation of MBone over Monterey BayNet for educational purposes. It documents the requirements and reconfiguration of Monterey BayNet sites to join the MBone. It shows that MBone over Frame Relay networks is possible and that current MBone technology provides great performance even on low speed network connections. This thesis shows how to configure and use frame relay networks which may be used by schools for distance learning.

5. Economical Storage and Retrieval of Digital Audio and Video for Distance Learning

Internetworking: Economical Storage and Retrieval of Digital Audio and Video for Distance Learning [Tiddy, 1996] is a thesis that focuses on the testing and comparison of currently available methods of digital A/V storage and the use of current transfer modes and protocols for on-demand retrieval of A/V over the Internet. Different compression methods, file formats, and World-wide Web applications will be examined. No new compression techniques are developed, but rather existing standards are being researched.

This thesis follows the "Recommendations for Future Work" section of an NPS thesis written by [Emswiler, 1995]. The Multicast Backbone (MBone) already provides the ability to send and receive near-real-time A/V over the Internet but currently provides no way to efficiently store the data and replay the session on demand. Tools that allow a user to play audio and video that is retrieved over the Internet do exist but force the user to transfer the entire file before it is played. The file sizes that accompany a 30 minute video, however, are so great (over 1 GB) that this method of retrieval is not a reasonable alternative due to data transmission speeds and user storage space limitations. Thus, practical compression and practical network distribution are the key bottlenecks to success.

6. Wide-area Network (WAN) for K-12 Schools

Internetworking: Planning and Implementing a Wide Area Network (WAN) for K-12 Schools [Bigelow, 1995] is a thesis that documents the planning, design, and implementation of a regional wide-area network connecting K-12 schools, research institutions, libraries, and institutions of higher education throughout the Monterey Bay area of California's central coast. The goal of the network is to enable students and educators to have access to the environmental information and resources available regionally via the Internet, at speeds which will encourage interaction and maintain interest. The wide-area network design process presents numerous human and technical challenges. These challenges are further amplified by the need to enable educators to design and implement school local area networks concurrent with the wide-area network

solution. The processes used to resolve these myriad issues and the resulting solutions are of direct relevance to the K-12 community as well as network planners, administrators and funding partners.

[Bigelow, 1995] details the problems and solutions in establishing a regional frame relay WAN which parallel many of the challenges encountered in establishing a regional ATM WAN in this thesis.

C. REGIONAL AND NATIONAL INTERESTS

1. BayNet and BayLink

Monterey BayNet is a regional WAN that connects K-12 schools, libraries, research institutions, and institutions of higher education throughout the Monterey Bay. Monterey BayNet is designed and implemented to increase the quality of education in the area. It enables students and teachers to access information and resources regionally via the Internet.

BayNet ATM is the Monterey Bay regional ATM network. BayLink is the ATM-based live lecture circuit between the Monterey Bay Aquarium, the San Jose Technical Museum of Innovation (SJTMI), and the Monterey Bay Aquarium Research Institute (MBARI) research vessel *Point Lobos*.

2. CalREN

In 1993 Pacific Bell (PacBell) created the California Research and Education Network (CalREN), a \$25 million statewide program to stimulate the development of new applications for high-speed data communications services. CalREN funds

collaborative projects whose applications revolutionize the ways organizations communicate and share information. CalREN application development work establishes a foundation for broadband services including state-of-the-art telemedicine research, diagnosis, and treatment; on-line schools with no geography, distance, or resource constraints; electronic democracy in the form of on-line, real-time interaction between government and citizens; and new business partnerships and ventures made possible by vast information storage, retrieval, and sharing capabilities.

NPS is involved in CalREN Project #ATMN-014: *ATM As An Enabling Technology for Tele-education, Telescience, and Electronic Libraries linking the Silicon Valley, Santa Cruz, and the Monterey Peninsula* [CalREN, 1994]. The project leader is Professor J. J. Garcia-Luna at the University of California, Santa Cruz (UCSC). The construction of the NPS ATM LAN has been heavily dependent on NPS and UCSC cooperation as well as CalREN support.

3. BAGNet

BAGNet is the first and largest of the CalREN projects and is the outgrowth of a several year effort to establish an ATM test bed in the San Francisco Bay area. The project participants include most of the academic, government, and industry computer science research organizations in the Bay area.

The goals of the project are to establish an ATM infrastructure, to engage all of the participants in an application that was uniquely suited to a high speed, metropolitan

area ATM network, and to enable several diverse applications that involved two or three of the participants directly collaborating. BAGNet is currently inactive.

4. I-WAY

The information wide area year (I-WAY) is an experimental high-performance network linking dozens of the country's fastest computers and advanced visualization environments. This network is based on ATM technology. The network supports both TCP/IP over ATM and direct ATM-oriented protocols. This network provides the wide-area, high-performance backbone for various experimental networking activities at Supercomputing '95 (SC'95). SC'95 took place December 3-8, 1995 at the San Diego Convention Center in San Diego, California. SC'95 was the premier conference for the presentation and discussion of research in high-performance computing and communications. SC'95 created a program that integrated fully with the capabilities of High Performance Computing and Communications (HPCC) and the Global Information Infrastructure (GII). The entire testbed was run on an ATM infrastructure.

It is built from a combination of existing network connectivity and some additional connectivity and services provided by multiple national service providers. The I-WAY is the most advanced infrastructure test bed to date to prototype the issues detailed in Figure 1.

- *Teraflop-class wide area computing:* Nodes of the network are top Supercomputing sites with a combined peak computing power approaching a teraflop of available computing. Much work is under way to make this distributed environment behave as one facility.
- *Close coupling of immersive virtual environments and Supercomputing:* Applications chosen for and developed over this network combine state-of-the-art interactive environments and back-end Supercomputing to tighten the link between computing and the user. The project is intended to demonstrate distance independence between resources, developers, and users.
- *An advanced application development resource:* The I-WAY is envisioned as a resource for advanced application development and demonstration. As such, much effort is being put into making it both user friendly and developer friendly.
- *Test bed to identify future network research issues:* The goal of this effort is to uncover the areas requiring further study and development. At a minimum, we are highlighting security mechanisms for wide-area computing, advanced end-to-end network management, the mapping of infrastructure to emerging application environments, and the mapping of applications to emerging infrastructure environments.

Figure 1. I-Way Prototyping Testbed, after [I-WAY, 1995].

5. IEEE Computer Graphics and Applications

In *Virtual Chesapeake Bay: Interacting with a Couple Physical/Biological Model*, [Wheless *et al.*, 1996] describe this multidisciplinary, collaborative project that fuses 3D visualizations of various types of data into a large-scale, interactive virtual world supporting investigation of coupled physical, biological, and environmental processes.

The Chesapeake Bay Virtual Environment (CBVE) provides a framework for integrating circulation and biological models with the computer visualization paradigm of the virtual world. The implementation of CBVE was demonstrated on the Global Information Infrastructure (GII) testbed constructed during SC'95. CBVE is a remote partner research application for the IIRG.

D. U.S. MILITARY INTERESTS

1. U.S. Navy

The U.S. Navy operates an ATM network at the Naval Research Laboratory (NRL). This network serves as a research and development test bed and supports application traffic as described in Figure 2.

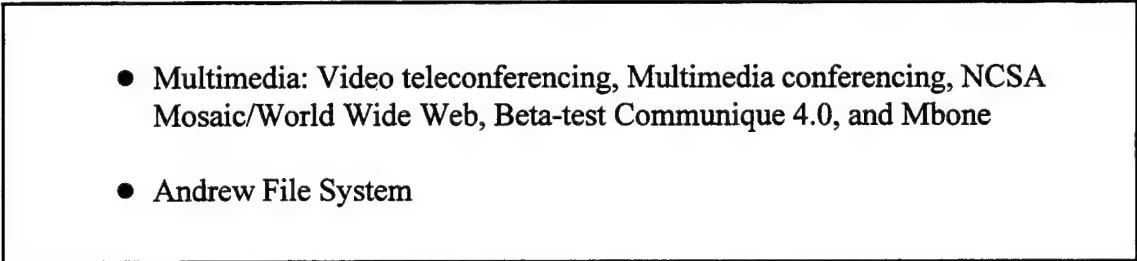
- 
- Multimedia: Video teleconferencing, Multimedia conferencing, NCSA Mosaic/World Wide Web, Beta-test Communique 4.0, and Mbone
 - Andrew File System

Figure 2. Applications of NRL's ATM Network, after [ATDNet, 1994].

The Center for Computational Science (CCS) conducts research as well as operational computation and network support at NRL. CCS supports basic and applied research at NRL and throughout the DoD. CCS has network connectivity to the SURANet National Science Foundation (NSF) regional network, the Defense Simulation Internet (DSI), the NASA science network, the Defense Research and Engineering

Network (DREN), and a regional ATM/FDDI/Ethernet network called the Advanced Technology Demonstration Network (ATDNet).

ATDNet is a high performance networking test bed in Washington, D.C. It is intended to be representative of possible future metropolitan area networks (MANs). It is established by the Advanced Research Projects Agency (ARPA) to enable collaboration among DoD and other federal agencies. ATDNet's primary goal is to serve as an experimental platform for diverse network research and demonstration initiatives, including deployment of ATM and Synchronous Optical Network (SONET) technologies [ATDNet, 1994].

NPS is examining the benefits and problems with ATM technology in order to provide NRL and other Navy labs with the feedback needed to correctly evaluate the strengths and weaknesses of ATM relevant to deployment afloat and ashore.

2. DoD

The Defense Information Systems Agency (DISA) has focused on ATM as being the network topology of the future. The following quote from DISA shows its commitment to ATM for the future of the Defense Information System Network (DISN):

[DISA] recognizes ATM technology as the key building block of the evolving DISN. It alone can satisfy the warfighter's need for the extension of high bandwidth, realtime and multi-media communications to remote theaters of operation anywhere in the world. It alone holds the promise of a single technology supporting high performance data, voice and video services through the seamless integration of local and wide area networks, including those in the tactical theater of operations.... Today ATM stands

alone with its promise of solving the vexing problems of seamless sensor-to-shooter information support to the warfighter. [DISA, 1994]
NPS is examining the benefits and problems with ATM technology in order to

provide DISA with the feedback it needs to correctly evaluate the risk of transferring to a cutting-edge technology such as ATM.

E. SUMMARY

This chapter discusses work related to the NPS ATM LAN. It first begins with local research occurring within the IIRG at NPS concerning network monitoring, network security, multicast and Mbone monitoring issues, a frame-relay WAN, and digital data storage and retrieval. Regional interests of BayNet, BayLink, CalREN, and BAGNet, and national interests such as the I-WAY and CBVE are examined. It concludes with military interests of the U.S. Navy and DoD. Much work has been done in this area and much work remains.

III. PROBLEM STATEMENT

The purpose of this research is to develop an Asynchronous Transport Mode (ATM) Local-Area Network (LAN) at the Naval Postgraduate School (NPS). The purpose of this LAN is to perform research into the capabilities of ATM and to enable various research projects using ATM technologies.

This thesis summarizes the hardware and software theory underlying ATM technology, describes how the LAN is constructed, provides experimental results obtained while building the NPS ATM LAN, and makes recommendations for the future use of ATM in the Navy.

Chapter II details research projects at NPS that relate to the NPS ATM LAN: ATM monitoring, security, ATM Mbone connectivity and monitoring, and digital data storage and retrieval. This chapter also discusses regional and national projects that interface with the NPS ATM LAN: BayNet/BayLink, CalREN, BAGNet, I-WAY, and CBVE.

Chapter IV delineates the details of the hardware methodology for the NPS ATM LAN — how the LAN is built. This chapter discusses the sequence of events taken in building the LAN by detailing the series of configurations that are used: stand-alone, peer-to-peer, single-switched, dual-switched, and finally connecting the LAN to the CalREN ATM WAN.

Chapter V specifies the software methodology behind ATM in general and the NPS ATM LAN in particular. It details ATM's peculiarities in comparison to traditional connectionless systems and provides a discussion of problems that exist when adapting internetworked (connectionless) multicasting to a (connection-oriented) ATM network.

Chapter VI documents the experimental results of the building and implementing of the NPS ATM LAN along the way. This chapter details the series of outcomes for operating stand-alone, peer-to-peer, single-switched, dual-switched, and finally connecting the LAN to the CalREN WAN. In particular, establishing multicast connectivity over the regional WAN is seen as the primary measure of effectiveness for internetworked ATM.

Chapter VII presents conclusions and recommendations for future work.

IV. HARDWARE METHODOLOGY

A. INTRODUCTION

As with most institutions, NPS has to deal with many years of legacy systems on a daily basis. [Leahy, 1988] and [Wiedenhoeft, 1994] exhaustively and exhaustingly present the evolution of the NPS computer network architecture. In his *Analysis of the Naval Postgraduate School Computer Network Architecture*, [Wiedenhoeft, 1994] recommends taking the campus backbone and collapsing it into an ATM WAN, as does Dave Norman [Norman, 1996].

This chapter details the present campus fiber backbone topology, how the NPS ATM LAN is constructed and tested, and the rudiments of what a future ATM network on the NPS campus might look like. Chapter VI details experimental test results, and Appendices D and E describe the setup of the Fore NICs and the ATM switches, respectively.

B. PRESENT CAMPUS FIBER TOPOLOGY

Figure 3 shows the campus' legacy fiber optic backbone topology. Located in Ingersoll Hall in the computer center (CC) was a Cisco A-100 HyperSwitch that connected to the California Regional Educational Network (CalREN) but did not terminate anywhere on campus. Fiber optic lines run from Ingersoll to Root Hall, from Ingersoll to Spanagel, and from Root to Spanagel. The systems technology lab (STL) is located in Root hall, and the Computer Science (CS) department is located in Spanagel hall.

NPS Campus Backbone Topology

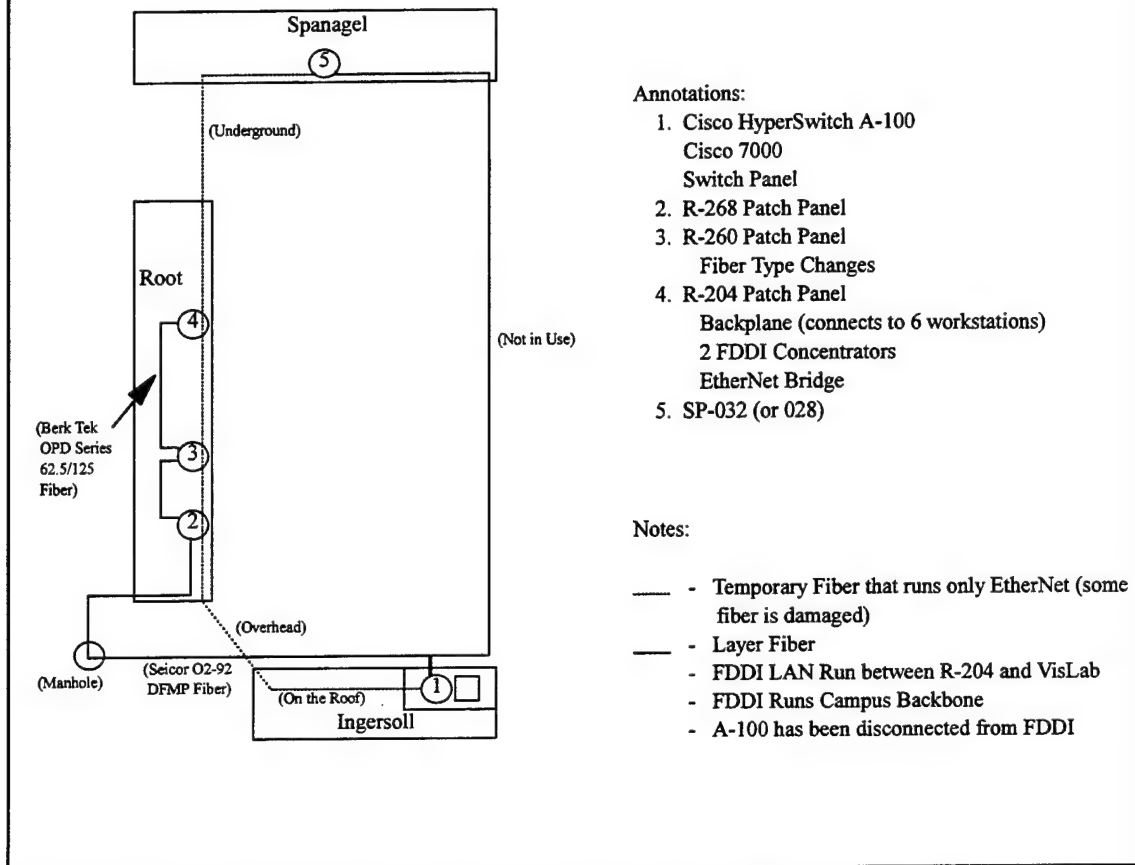


Figure 3. NPS Campus Backbone Topology, after [Romo, 1995].

The installed fiber optic backbone was not previously used for ATM. However, FDDI and Ethernet are running from Ingersoll to both Root and Spanagel. The line running from Ingersoll to Spanagel is unused (i.e. "dark" fiber).

C. DESIGNING THE FUTURE ATM NETWORK

I requested a meeting of the NPS networking advisory group to discuss the future of ATM at NPS and how to build the campus area network (CAN). The following personnel were present at the November 22, 1995 meeting: Dale Courtney (student), Rob Williams (student), Don Brutzman (Assistant Professor), Mike McCann (CC), Terry

Williams (STL), Mark Will (CS), Sue Whalen (CS), and David Pratt (CS). The discussion first focused on available campus ATM resources as shown in Figure 4. Second, the activation priorities for the ATM LAN were decided upon. The planned sequence of events is outlined in Figure 5. Third, a decision is reached as to what is needed to reach the goal, as listed in Figure 6.

Following this meeting, I further discussed the NPS ATM LAN with the NPS network administrator, Roy Romo (CC). Figure 7 details his concerns about building the ATM LAN. His concerns are soon put to rest, as Figure 8 describes. After solutions are found addressing the computer center's concerns, the actual task of constructing the NPS ATM LAN began.

- Three Cisco A-100 switches (one installed in CC)
- One Cisco 7000 switch (also installed in CC)
- Two ATM NICs for an Indy workstation
- Two ATM NICs for a Challenge workstation
- Ten ATM cards for the A-100 switches
- Many workstation candidates for dual-homing (ATM/IP)

Figure 4. ATM resources, after [Advisory Group, 1995].

- STL ATM LAN — get a switch working in Rt-204
- VisLab Hookup — get ATM between Rt-204 and the VisLab
- Other hookups
 - CalREN Connection
 - Spanagel/CS
 - Professor Nuss, Root-245

Figure 5. Activation Priorities, after [Advisory Group, 1995].

- Verified the status of fiber lines running to Spanagel, ensuring that the lines presently installed are sufficient to support ATM in the future.
- Ascertained from Professor Nuss:
 - That he has a workstation.
 - What he needs in the way of equipment (he is working on a joint oceanographic project with UCSC).
 - That he needs a NIC. UCSC is providing this.
 - That Russ Schwanz is the System Administrator for Professor Nuss' system.
 - Two fiber lines need to be routed to his office from the Rt-204 patch panel.
- Got written permission for Terry Williams (STL) to hook up to CalREN
- Got a commitment from 05 that the topology between CC, STL, and CS will be circular vice star
- Got a work definition request from CS to Terry Gentry (CC) that specifies what CS needs from CC to hook up ATM service to Spanagel
- Requested IP addresses for our ATM LAN
- Verified the type of NICs that the VisLab needs

Figure 6. Requirements to Meet the Goal, after [Advisory Group, 1995].

- Who funds the fiber that needs to go in place?
- The fiber in place (62.5 x 125 standard multimode) probably will not do. Single Mode (SM) is required.
- Fiber connectors in place are not compatible with SM.
- How much more physically fragile is SM over Multi Mode (MM) cables?
- Need to test out SM lines. OC-3 requires SM fiber.
- Whether the LAN will use IP over ATM or LAN emulation
- Who buys the NICs?
- Building this LAN cannot be done cheaply with what is in place.

Figure 7. Network Administrator Roy Romo's Concerns, after [Romo, 1995].

D. CHRONOLOGY OF EVENTS

What follows is the sequence of events detailing how the ATM LAN is built on campus. The ease with which the ATM LAN came together is due to system administrator expertise and the NPS networking advisory group's agreeing to the sequence of events before construction began. The details of configuring the Fore cards are contained in Appendix D, configuring the Cisco A-100 switch in Appendix E, and experimental test results in Chapter VI.

- There will be no need to lay any more fiber since the legacy system between the three major computing areas (CC, STL, CS) will work well.
- SM fiber is not required except for long distances or extremely high data rates. The CAN supports short distances, and the A-100 switches can only provide up to OC-3 (155 Mbps) and are designed to work with MM fiber [Cisco, 1995].
- The incompatibility of connectors is easily fixed with the use of adapters. The SM fiber is no more physically fragile than MM fiber.
- No reason to wait for SM fiber to be laid since MM works well at OC-3 rates.
- For ease of use in connecting to BayNet, the LAN is designed to run IP over ATM.
- The respective departments (STL, CS, CC) will purchase the NICs for their own users so that CC will not have to pay for the entire campus' development costs.

Figure 8. Putting the Concerns to Rest.

1. Stand Alone

The first step in the installation process is to select two Silicon Graphics workstations in STL on which to build the initial portion of the ATM LAN. One is the Indy Workstation "Royal" at Ethernet IP address 131.120.63.16, and the other the Challenge DM workstation "Navy" at Ethernet IP address 131.120.64.23. Figure 9 shows these two workstations in the stand alone configuration.

An ATM NIC is installed and configured in each of these workstations. The specifications for the Fore NICs are detailed in Appendix B. The configuration details for

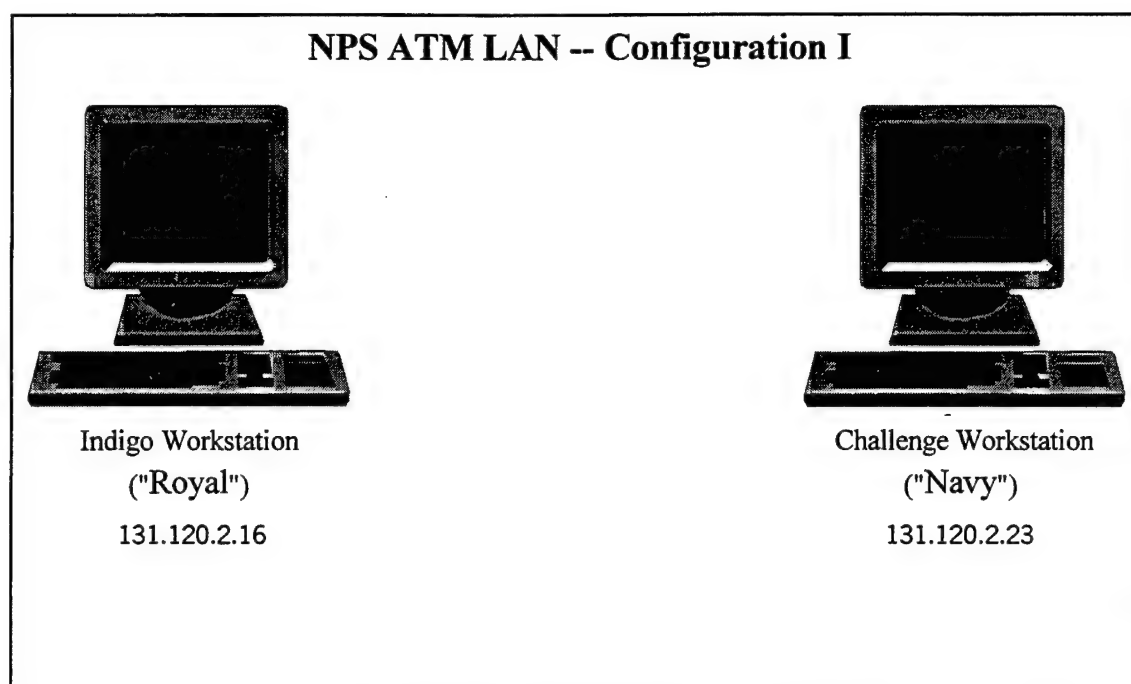


Figure 9. Stand Alone Configuration.

the Fore NICs are detailed in Appendix D.

Royal's ATM address is set as 131.120.2.16, borrowing from STL's 2-Net until CC assigns permanent IP addresses; Navy is set as 131.120.2.23. Royal had no room for the ATM NIC, so the FDDI interface is removed from Royal and replaced with the Fore ESA-200 ATM NIC. The second VME ATM card is given to Mike McCann (CC) in the VisLab for his Power Onyx. The Power Onyx also had no free slots in it, so ATM is not being installed in the VisLab until CC can resolve these hardware problems.

The Fore software is loaded in the */usr/etc/fore/etc* directory on both Navy and Royal. "Route back" testing (described in Chapter VI) is performed on each machine to verify that Fore's ATM hardware and software are working correctly on each workstation. Testing results are satisfactory, and Chapter VI describes the experimental

results.

2. Hooking up Two Workstations Peer-to-Peer

The second step in the installation process was to connect the two workstations with a single pair of fiber optic lines without using a switch. Figure 10 shows the two workstations in this configuration.

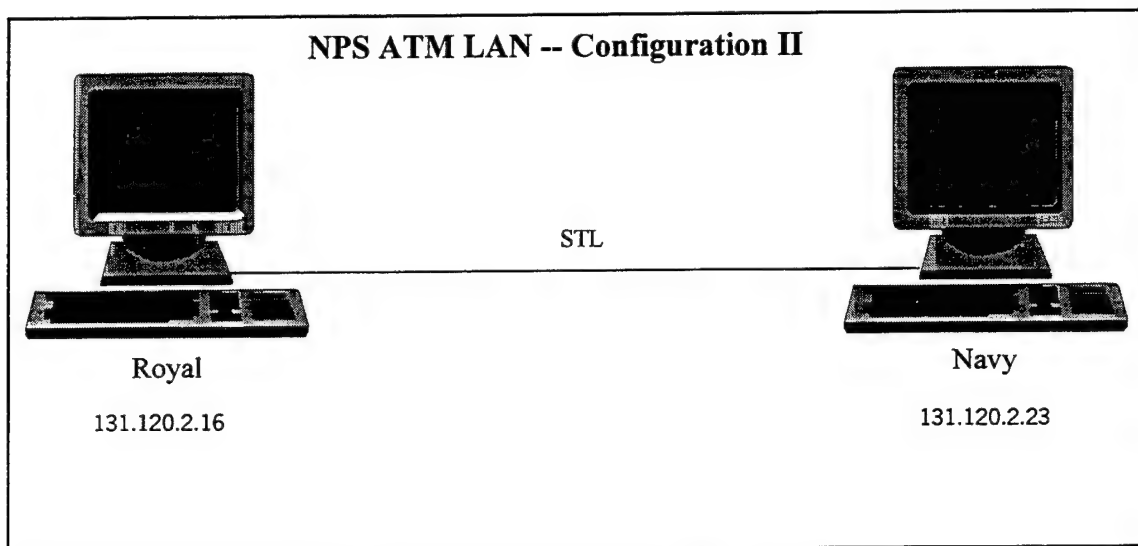


Figure 10. Peer-to-Peer Configuration.

Initially this configuration did not work. Troubleshooting the constantly red receive light on the ESA-200 card reveals that the fiber optic cable is not connected correctly — the presumed Transmit (T)-Receive (R) configuration is actually R-T. After swapping the cables, cells are successfully passed between the two machines by using the Unix ping, ftp, telnet, rlogin commands, and routeback testing. Chapter VI describes the results of these tests. Figure 11 displays what the dual-homed workstations look like when attached to a single fiber cable without a switch.

3. Hooking up a Single ATM Switch to Two Workstations

The third step in the installation process is to connect the two workstations with a single Cisco ATM A-100 HyperSwitch. Figure 12 shows the two workstations in this configuration. The specifications for the switch are detailed in Appendix C. The configuration of the switch is described in Appendix E. Again, cells are successfully passed between the two machines by using the Unix ping, ftp, telnet, rlogin commands and routeback testing. Chapter VI describes the results of these tests.

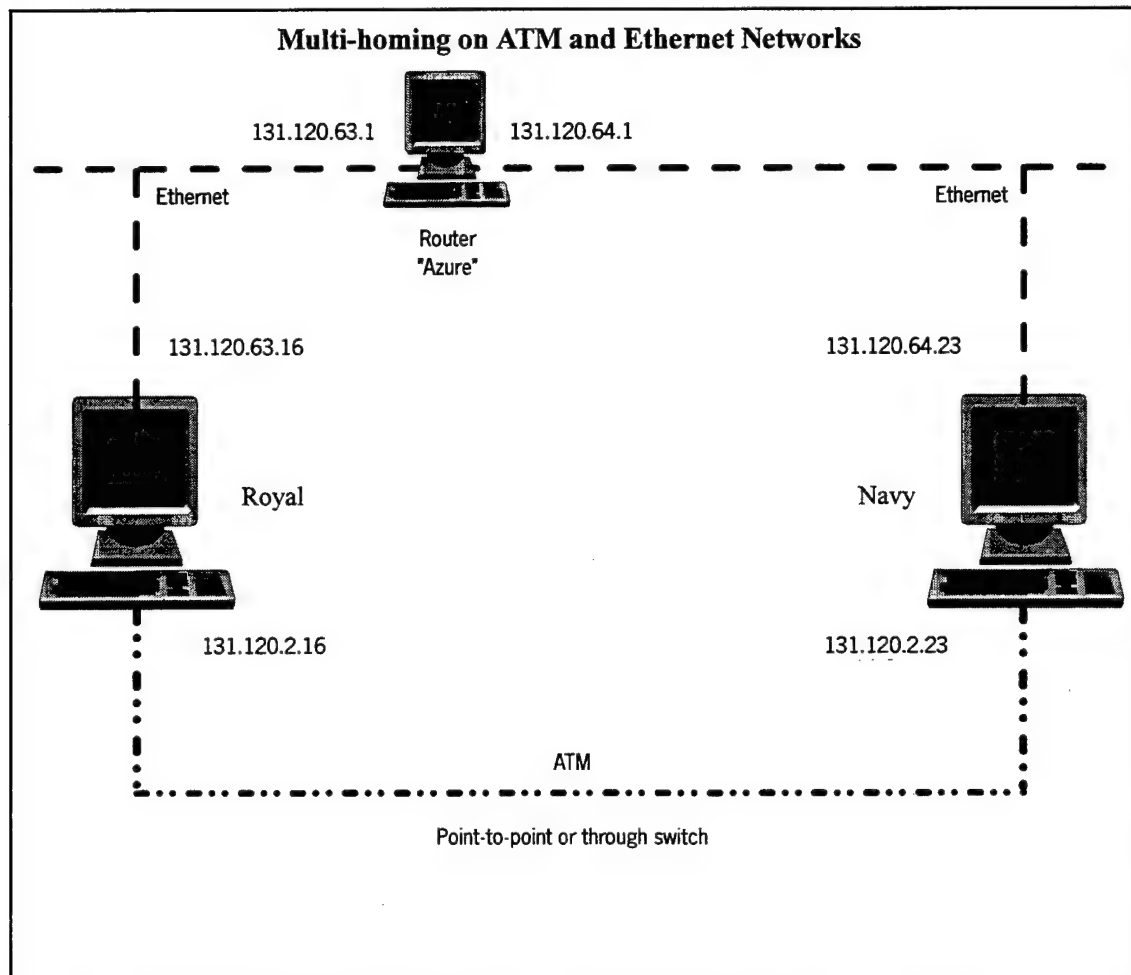


Figure 11. Multi-homing of ATM and Ethernet.

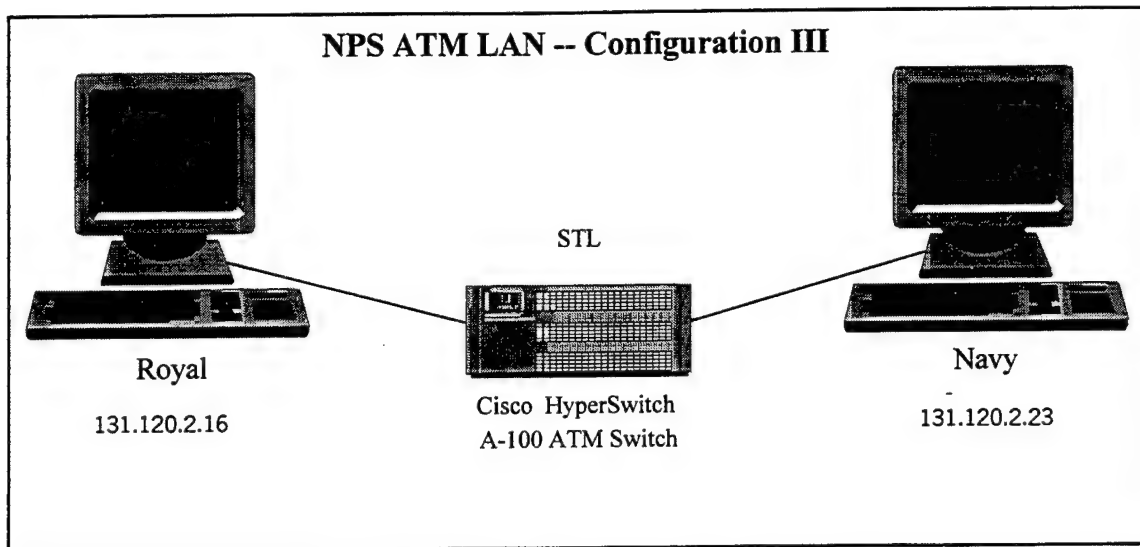


Figure 12. Single-Switched ATM LAN.

4. Hooking up Two ATM Switches

The fourth step in installation process is to connect the two workstations with two switches, as shown in Figure 13. Cisco ATM A-100 HyperSwitches are used for both switches. The configuration of the switches is described in Appendix E.

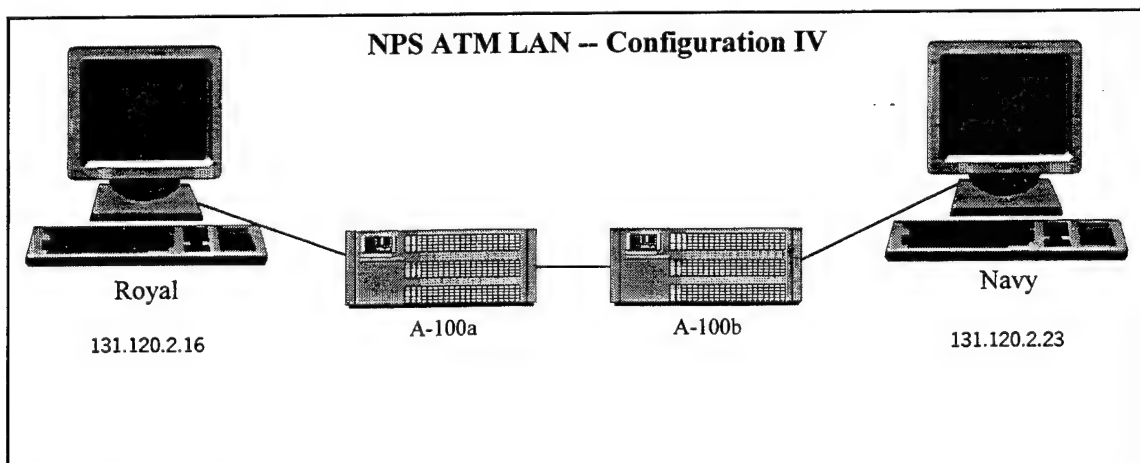


Figure 13. Dual-Switched ATM LAN.

Cells are again successfully passed between the two machines by using the Unix ping, ftp, telnet, rlogin commands, and routeback testing. Chapter VI describes the results of these tests.

5. Connecting to BayNet

The final step in the installation process is to connect the school's ATM LAN to CalREN and BayNet. Figure 14 displays the initial configuration.

Our immediate goal is to communicate with UCSC over the BayNet cloud. Details of the BayNet cloud are shown in Figure 15.

The first thing that has to be accomplished is connecting the ATM LAN in Root hall to the computer center's A-100 switch. That switch is presently connected to the

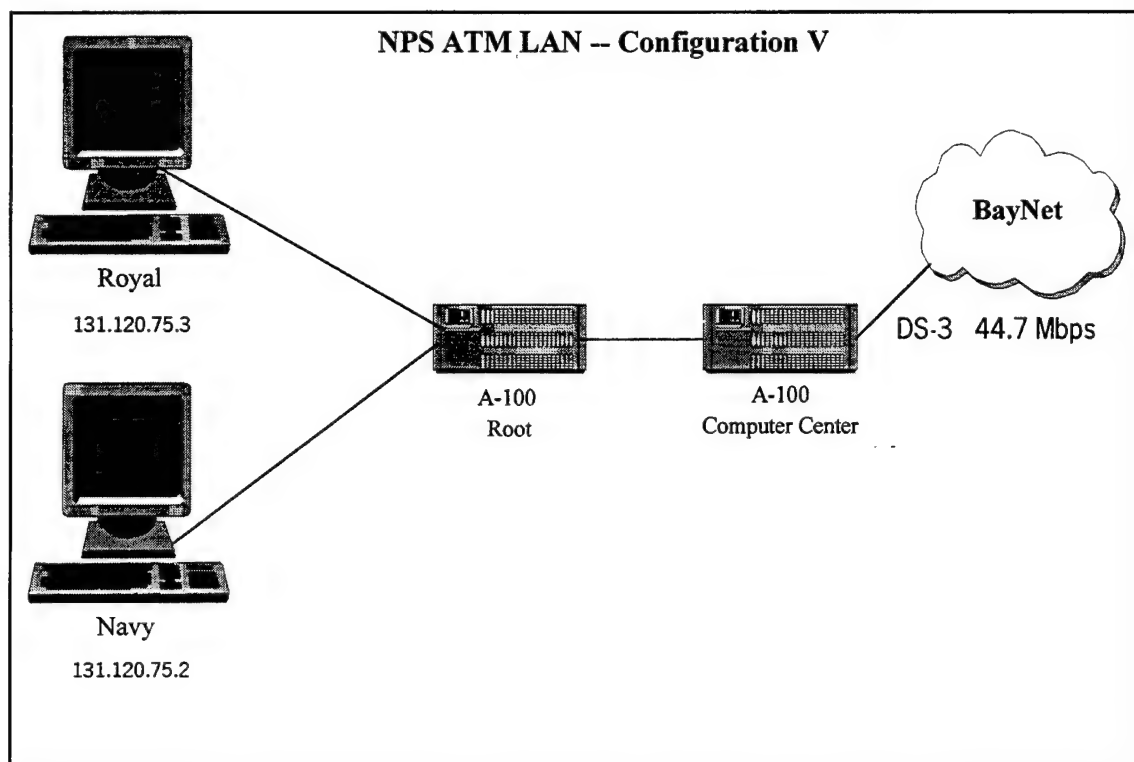


Figure 14. NPS ATM LAN Connection to BayNet Cloud.

CalREN network with a DS-3 (44.7 Mbps) connection, but that circuit does not continue past the CC switch into NPS. The configuration of the switches is described in Appendix E. CC assigns a subnet address for the ATM LAN of 131.120.75.0. Navy is assigned 131.120.75.2 and Royal 131.120.75.3.

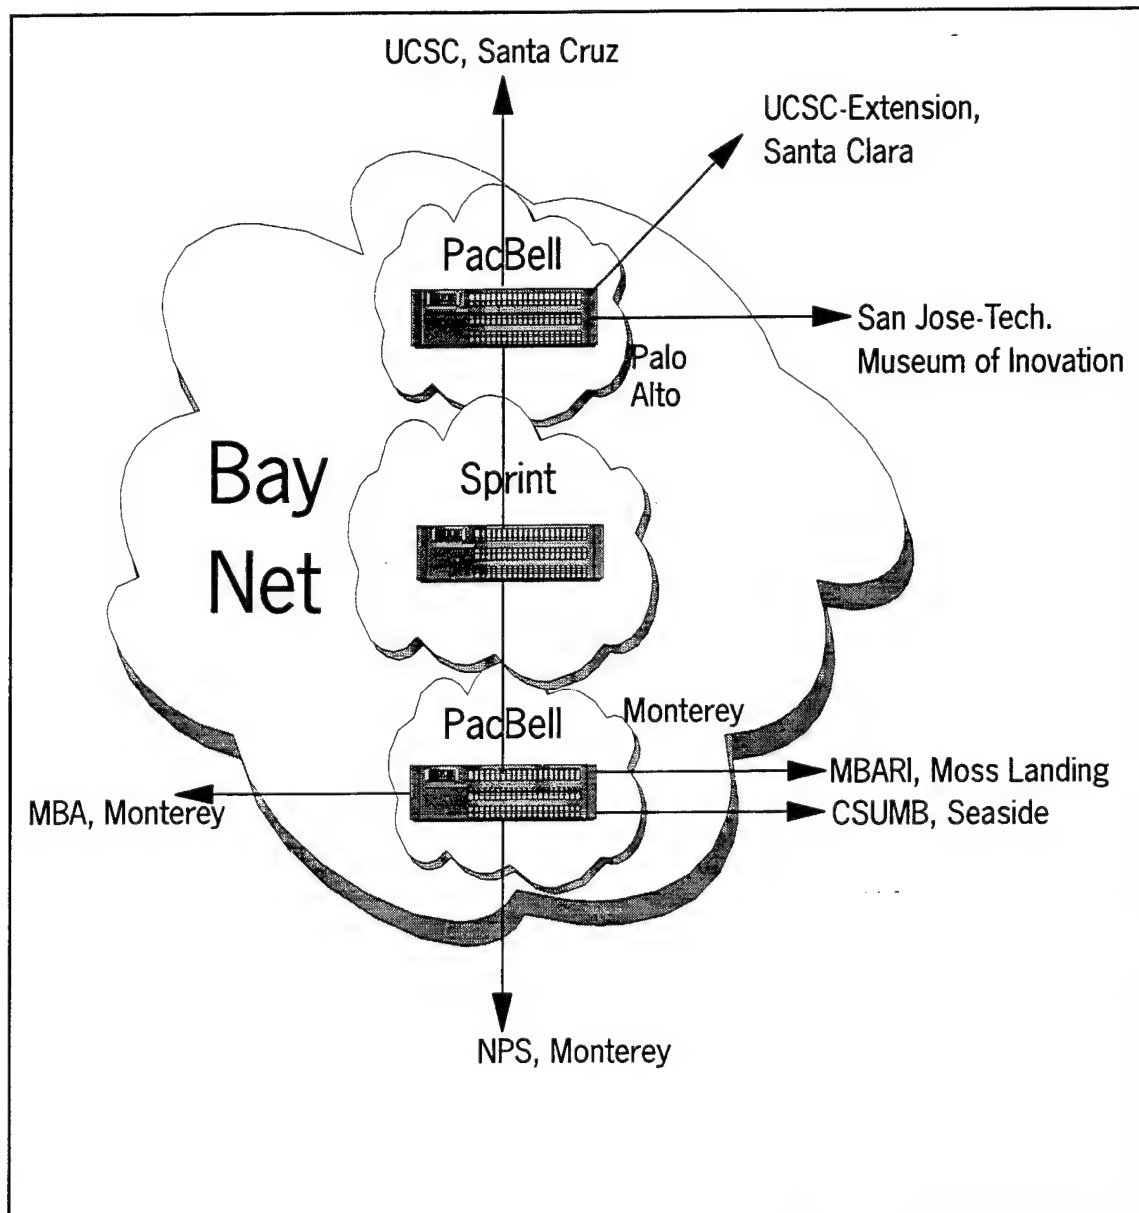


Figure 15. BayNet Configuration, after [Arul, 1996].

UCSC agrees to perform ATM ARP functions on their ASX-200 switch until the NPS Cisco 7000 switch can be configured to perform that function. The Monterey Bay Aquarium (MBA) uses ATM to communicate exclusively with the San Jose Technological Museum of Innovation (SJTMI) via PVCs. The Monterey Bay Aquarium Research Institute (MBARI) and California State University at Monterey Bay (CSUMB) are not online with ATM yet.

Many problems exist in establishing connectivity with UCSC. UCSC is using exclusively Fore ASX-200 switches and Fore NICs for its connectivity both on campus and to its off-campus extension. UCSC uses SVCs to run all its multimedia applications across the Pacific Bell (PacBell) Palo Alto (PB-PA) switch. Figure 17 displays the

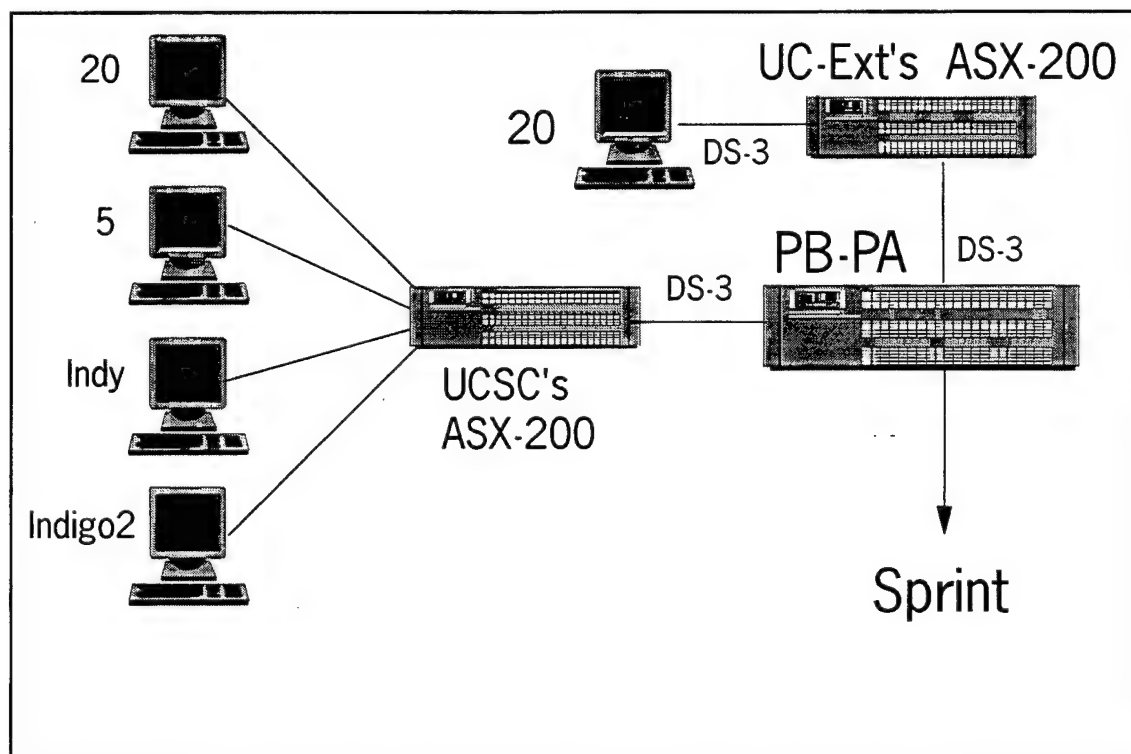


Figure 16. UCSC's ATM Configuration, after [Arul, 1996].

physical connections for UCSC.

NPS has been unable to get an SVC to UCSC since the Sprint switch is unable to support SVCs. As Figure 17 shows, all traffic between the upper local access and

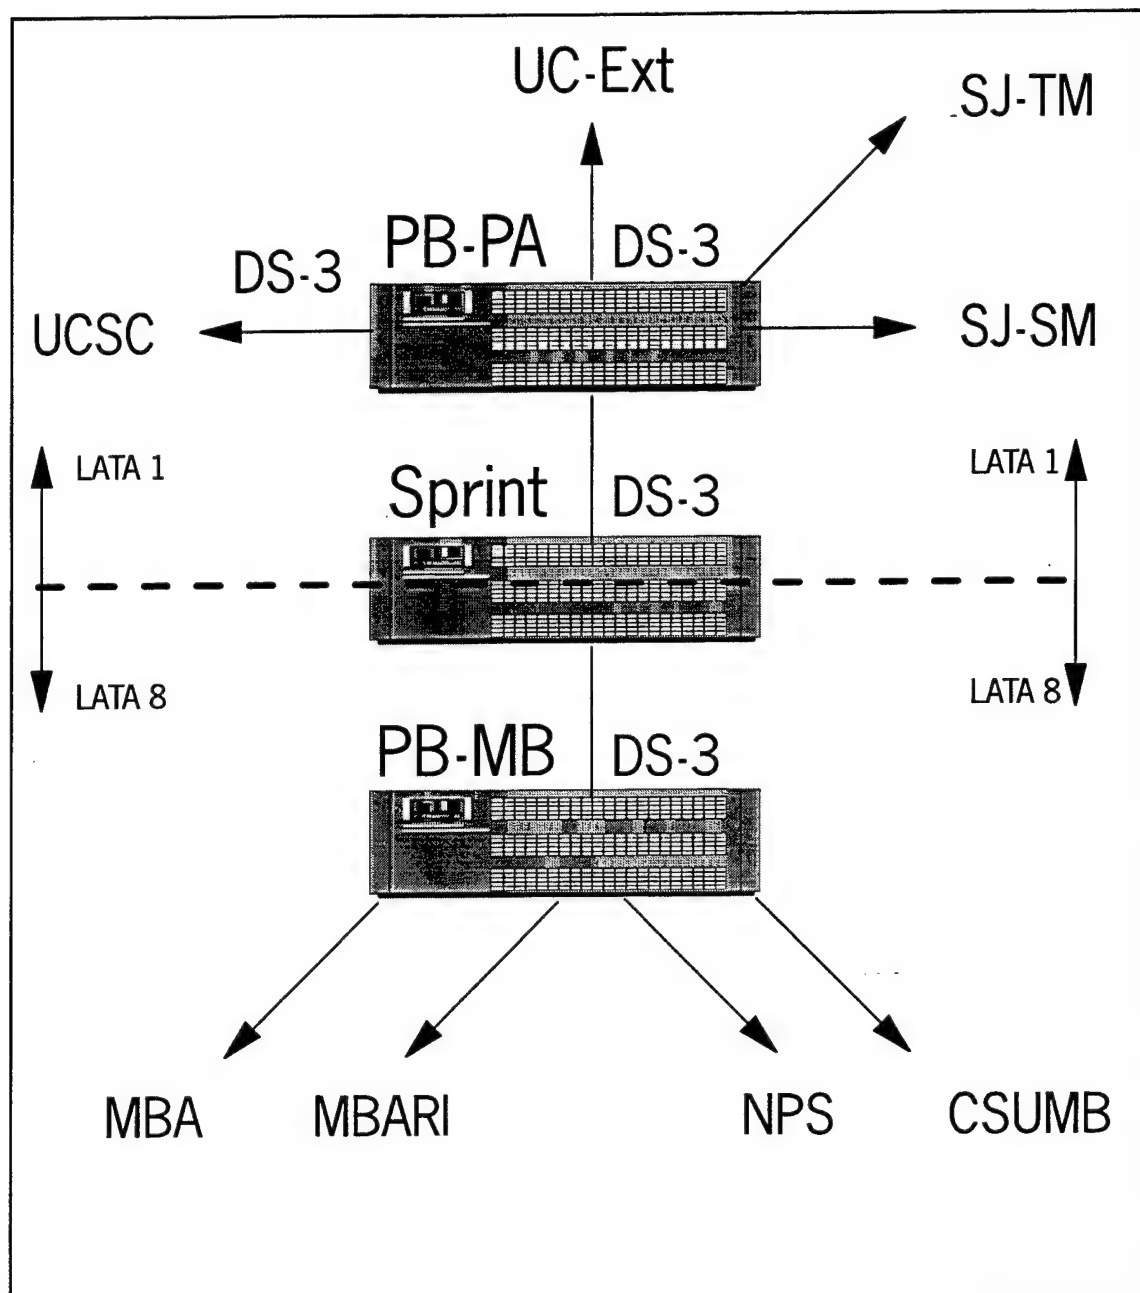


Figure 17. Sprint's and PacBell's Switches, after [Arul, 1996].

transport area (LATA 1) and the lower service area (LATA 8) is required to go through a long-haul carrier — in this case Sprint — prior to the telecommunications act of 1996 (96 Act) . This artificial regulatory artifact inhibits interoperability by unnecessarily complicating regional connectivity.

However, on February 8, 1996, the President signed into law the long-awaited Telecommunications Act of 1996. This legislation represents the first major overhaul of the Communications Act of 1934. The 96 Act removes statutory, antitrust, and other restrictions on telecommunications carriers and others who may now compete and combine more or less freely with one another in telephone, cable, and television markets. The legislation will also affect business, educational, and residential customers who will experience fundamental changes in the manner in which they purchase and receive telecommunications services. [Sapronov, 1996]

The Bell operating company (BOC) may now offer out-of-region and incidental interLATA services immediately upon enactment of the 96 Act. InterLATA provisions of the 96 Act are listed in Figure 18 [Sapronov, 1996]. When PacBell begins to provide all of the ATM network services, the interoperability problems experienced now due to Sprint's inability to support SVCs should be resolved. Curiously, despite FCC rules issued in August 1996, PacBell managers indicate that an additional year may be required before interLATA charges might be eliminated.

- Audio or video programming which is capable of interaction with subscribers
- Alarm monitoring services
- Two-way interactive video services or Internet services
- Commercial mobile services
- A service that permits a customer that is located in one LATA to retrieve stored information from a database storage in another LATA
- Signaling information used in connection with the provision of telephone exchange services
- Network control signaling information to common carriers offering interLATA services

Figure 18. InterLATA provisions of the 96 Act , after [Sapronov, 1996].

Since the Sprint switch is unable to support SVCs, an agreement is reached between PacBell and Sprint to provide a 10 Mbps PVC for communications between NPS and UCSC. UCSC assigns a subnet address of 172.20.70.0 to the NPS ATM LAN. Subsequently, Royal is assigned the address 172.20.70.1, and UCSC's "Cyclone" workstation is assigned 172.20.70.2. That PVC is tested satisfactorily by using the Unix ping, ftp, telnet, rlogin commands, and routeback testing. The experimental results are documented in Chapter VI. The configuration of the NPS LAN is now complete, and the final overview is displayed in Figure 19.

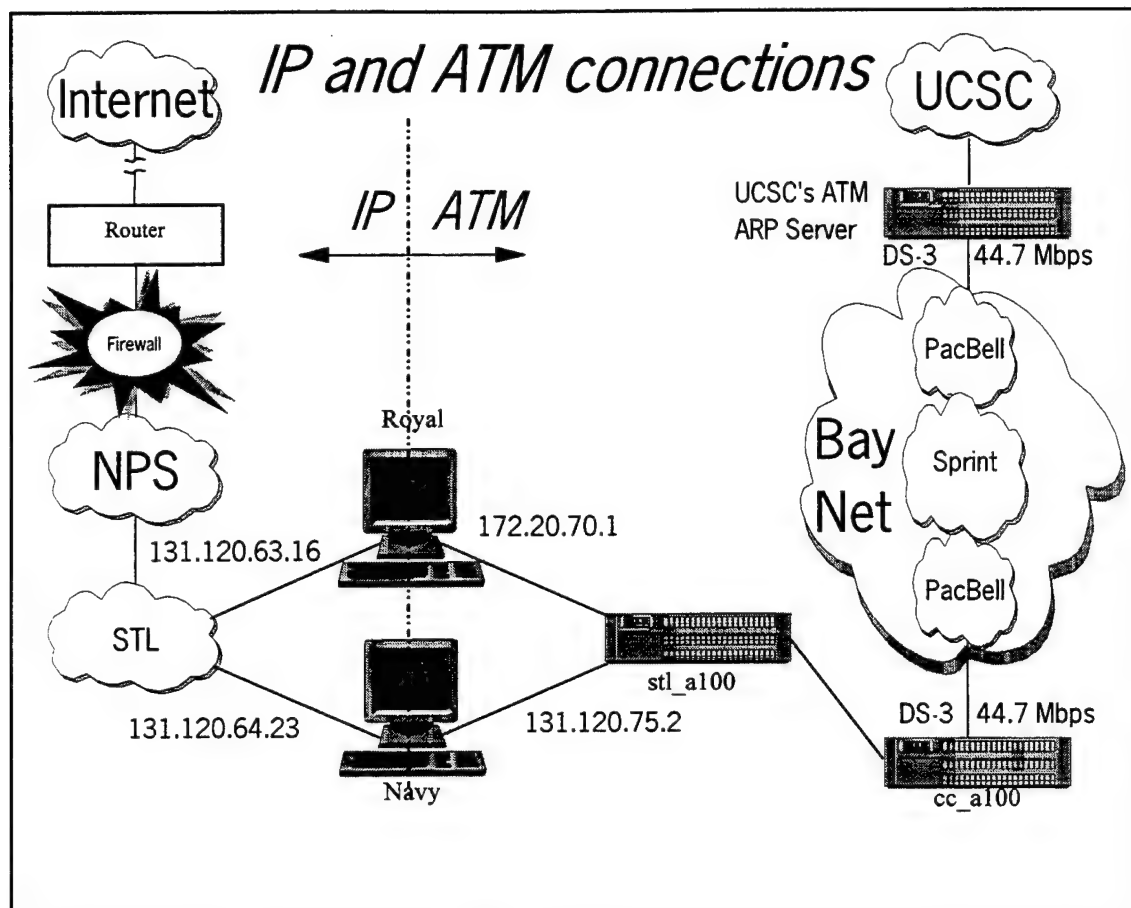


Figure 19. NPS ATM LAN — IP and ATM connections.

E. FUTURE DESIRED CONFIGURATION

The future desired configuration is shown in Figure 20. This reflects the short-term backbone network for the campus that connects Ingersoll (VisLab and CC), Root (STL and Professor Nuss), and Spanagel (CS) halls with BayNet. All of the internal lines will be running at OC-3 (155 Mbps). The BayNet connection is limited to DS-3 (44.7 Mbps) by PacBell.

F. SUMMARY

The physical construction of the NPS ATM LAN was completed easily. This was mostly due to highly qualified network administrators and the strong, up-front design considerations developed by the NPS networking advisory group. All questions and unknowns were resolved prior to implementation commencing.

The five-step process is to install ATM in a stand-alone configuration, then peer-to-peer, through one switch, through two switches, and finally across campus and out to

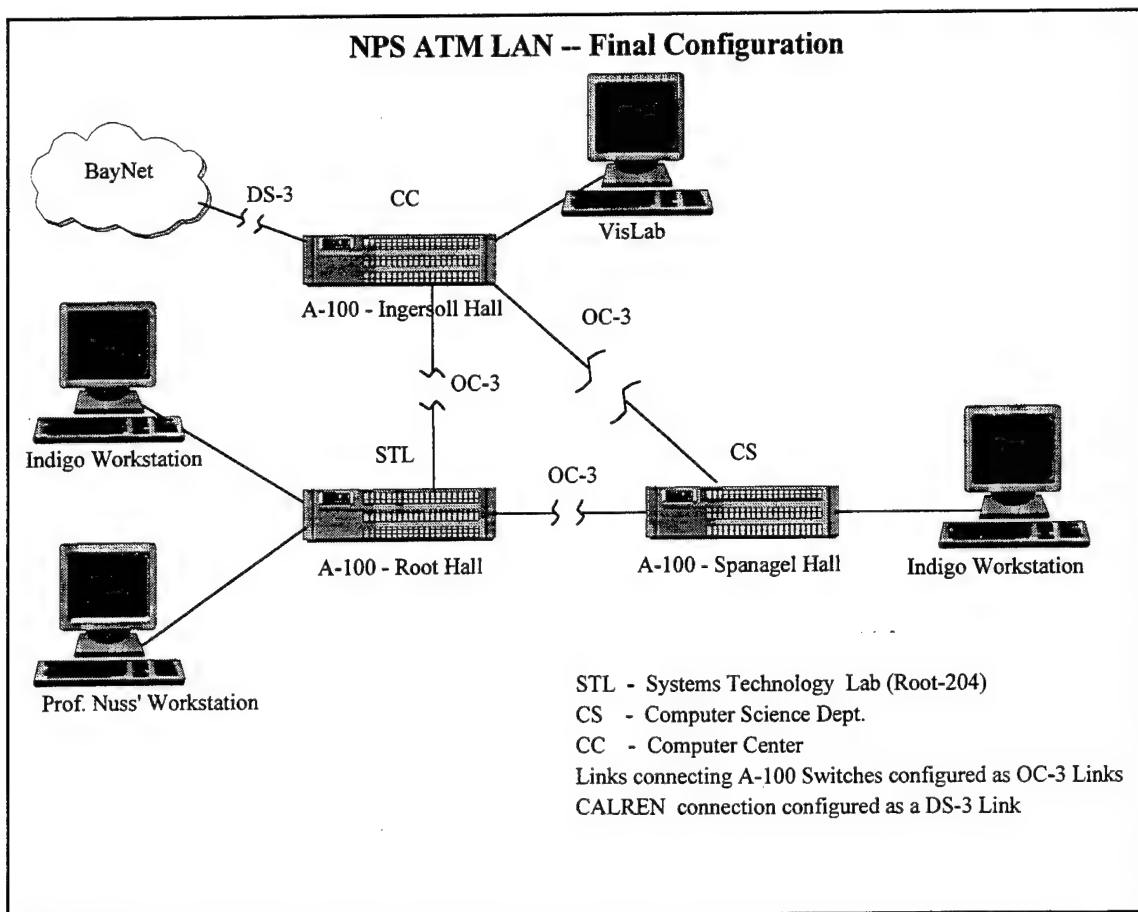


Figure 20. Final Desired Configuration, after [Advisory Group, 1995].

BayNet. The NPS ATM LAN of the future will also connect Spanagel (CS), Professor Nuss (Oceanography), and the VisLab (CC), providing OC-3 data rates to all major computing hubs across campus with an ATM backbone.

V. SOFTWARE METHODOLOGY

A. INTRODUCTION

This chapter develops the theoretical background of ATM in general and then applies it specifically to the NPS ATM LAN. Related sections include Appendix D, which discusses the specifics of configuring the Fore Network Interface Cards, and Appendix E, which discusses the details of configuring the Cisco A-100 ATM switches.

B. SIMPLE ATM SUMMARY

[Partridge, 1994] outlines four major assumptions behind the design of ATM. These assumptions are listed in Figure 21.

ATM is a protocol that transmits data as fixed sized packets — one culmination of many developments in switching and transmission of data in the last twenty years. It was designed to make Broadband-Integrated Services Digital Network (B-ISDN) a reality. B-ISDN was created conceptually as just an extension of ISDN, so it functions as a

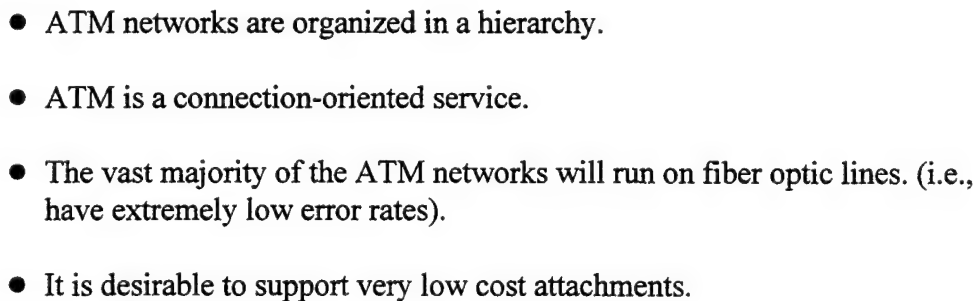
- 
- ATM networks are organized in a hierarchy.
 - ATM is a connection-oriented service.
 - The vast majority of the ATM networks will run on fiber optic lines. (i.e., have extremely low error rates).
 - It is desirable to support very low cost attachments.

Figure 21. Four Major Assumptions Behind ATM, after [Partridge, 1994].

communication network that can provide integrated broadband services.

The basic operation of an ATM switch is simple: the switch receives a cell across a link, looks up the connection value in a local translation table to determine the outgoing port (or ports) of the connection, and then retransmits the cell on that outgoing link with the appropriate connection identifiers. [Partridge, 1994] Thus, transmission of cells along established links is very efficient. However, setting up appropriate links can be very complex.

C. ATM PROTOCOL REFERENCE MODEL

ATM architecture is organized into layers and planes. Layers deal with distinct levels of capabilities or services that build upon each other. Communication from higher layers is adapted to lower ATM layers which in turn pass the information on to the physical layer for transmission over a selected physical medium [Siu, 1994]. The three layers created in the ATM hierarchy are listed in Figure 22. Planes is an ATM term which specifies domains of activity. The domains of activity distinguished for ATM are listed in Figure 23.

- The *control plane*, on which calls and connections are established and maintained.
- The *user plane*, on which users or nodes exchange data. This is the plane at which ordinary user services are provided.
- The *management plane*, on which network-management and layer-management services are provided. This plane coordinates the three planes and manages resources for the layers.

Figure 22. Three Domains (Planes) of ATM Activity, after [Feibel, 1995].

- *Physical* — Defines the bit timing and other characteristics for encoding and decoding the data into suitable electrical/optical waveforms for transmission and reception on the specific physical media used. It also provides cell delineation function, header error check (HEC) generation and processing, performance monitoring, and payload rate matching.
- *ATM* — Cell headers and trailers are created, virtual channels and paths are defined and given unique identifiers, and cells are multiplexed or demultiplexed. The ATM layer creates the cells and uses the physical layer to transmit them.
- *Adaptive* — Interfaces the higher layer protocols to the ATM layer and relays ATM cells both from the upper layers to the ATM layer and vice versa. When relaying information received from the higher layers to the ATM layer, the AAL segments the data into ATM cells.

Figure 23. Three layers created in the ATM Hierarchy, after [Feibel, 1995].

1. The Physical Layer

In the mid-1980s, Bellcore proposed synchronous optical network (SONET) as the standard for optical transmission for the U.S. telephone system. At that time, there were no existing standards for transmission equipment at the physical or the framing level. SONET was an attempt to overcome this problem by defining framing standards [Yao, 1994], and is the ANSI standard [Feibel, 1995].

The ANSI-sponsored subcommittee T1S1 standardized SONET as the preferred physical layer for ATM. The SONET physical layer specifications establish a world-wide digital telecommunications network hierarchy [Ebrahim, 1992]. The CCITT counterpart to SONET is Synchronous Digital Hierarchy (SDH), which defines the STM-x channel capacities.

SONET standardizes transmission around the bit rate of 51.84 Mbps (STS-1), and multiples of this bit rate comprise higher bit-rate streams. These rates are based on legacy U.S. voice-service telephone systems running at 125 μ sec intervals (8 kHz) between synchronous frames to meet the Nyquist value for acceptable human voice quality [Freeman, 1991]. OC-3 is of particular interest as this is the lowest bit rate initially designed to carry the ATM traffic [Ebrahim, 1992].

SONET and SDH are similar but different enough that they do not interoperate. The major difference is that SONET is based on STS-1 at 51.84 Mbps for efficient carrying of T3 signals. SDH is based on STM-1 at 155.52 Mb/s for efficient carrying of E4 signals. The way payloads map into these building blocks is also different. From a

formatting perspective, OC-3/STS-3 \neq STM-1 even though the data rate is the same.

However, SONET STS-3c (STS-3 concatenated) is the same as SDH STM-1, and STS-9c = STM-3c, etc. Figure 24 shows the relationship between SONET, STS/OC, SDH, and STM. [Symborski, 1995]

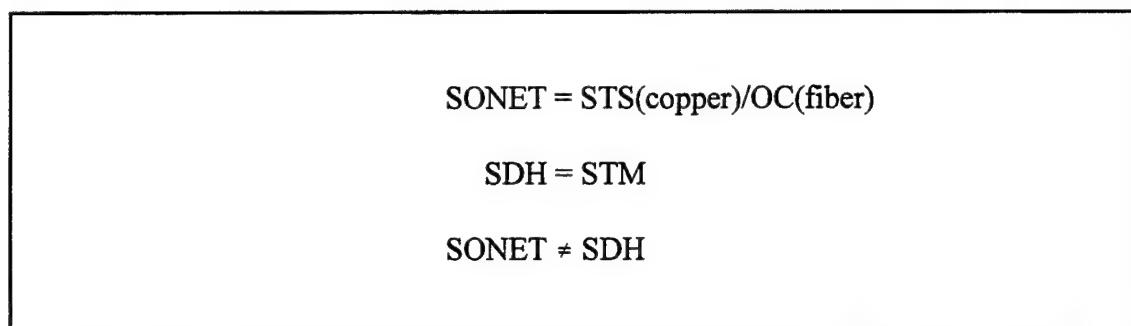


Figure 24. Relationship between SONET and SDH, from [Symborski, 1995].

The NPS ATM LAN runs at a maximum speed of OC-3/STS-3, and the PacBell connection between NPS and the CalREN Monterey Bay ATM WAN is DS-3. Table 1 shows standardized asynchronous data rates, and Table 2 shows the relationship between OC, STS, and STM.

ATM cells are transported across the physical layer as an opaque payload, also called the SONET payload or the synchronous payload envelope (SPE). Opaque implies that cell contents are not examined during switching. The physical layer is independent of the payload type and can just as easily carry STM cells as ATM cells. [Ebrahim, 1992]

Asynchronous Data Rates			
North America Designation	Europe Designation	Data Rate (Mbps)	Comments
DS0	-	64 kbps	1 Voice Channel (not compressed)
DS1 = T1	-	1.544	24 DS0s
-	E1	2.048	32 Voice Channels
DS1c	-	3.192	2 DS1s concatenated, indivisible payload
DS2	-	6.312	4 DS1s
-	E2	8.448	4 E1s
-	E3	34.368	16 E1s
DS3		44.736	28 DS1s, 7 DS2s

Table 1. Asynchronous Data Rates, after [Freeman, 1991; Partridge, 1994].

The physical layer has two sublayers. The lower sublayer, called physical medium (PM), includes the definition for the medium and the bit-timing capabilities which are all media dependent. The upper sublayer, called transmission convergence (TC), is responsible for making sure that valid cells are being created and transmitted and is media independent. TC breaks off individual cells from the data stream of the higher ATM layer, checks the cell's header, and encodes the bit values [Feibel, 1995]. This structure is the same as FDDI and 802.x LAN standards. Figure 25 shows the ATM physical layer.

SONET/SDH Transmission Rates		
SONET electric/optic STS-n/OC-n	Data Rate Mbps	Comments
		Note: STS-n could as well be OC-n
1	51.84	28 DS1 or 1 DSS
3	155.52	3 STS1 byte-interleaved
3c	155.52	3 STS1 concatenated, indivisible payload
9	466.59	9 STS1, 3 STS3c, or any mix
12	622.08	12 STS1, 4 STS3c, or any mix
12c	622.08	12 STS1 concatenated, indivisible payload
18	933.12	18 STS1, 6 STS3c, or any mix
24	1,244.16	24 STS1, 8 STS3c, or any mix
36	1,866.24	36 STS1, 12 STS3c, or any mix
48	2,488.32	48 STS1, 16 STS3c, or any mix

Table 2. SONET/SDH Transmission Rates, after [Freeman, 1991; Partridge, 1994].

The ATM Forum allows for various types of physical interfaces for ATM networks, including the fiber optic interfaces listed in Figure 26.

As discussed in Chapter IV, the NPS ATM LAN's internal interfaces are all OC-3. The connection to Monterey BayNet is DS-3.

Over long distances, such as within the public switched telephone network (PSTN), the cells are encapsulated inside SONET frames. There are rules for how ATM

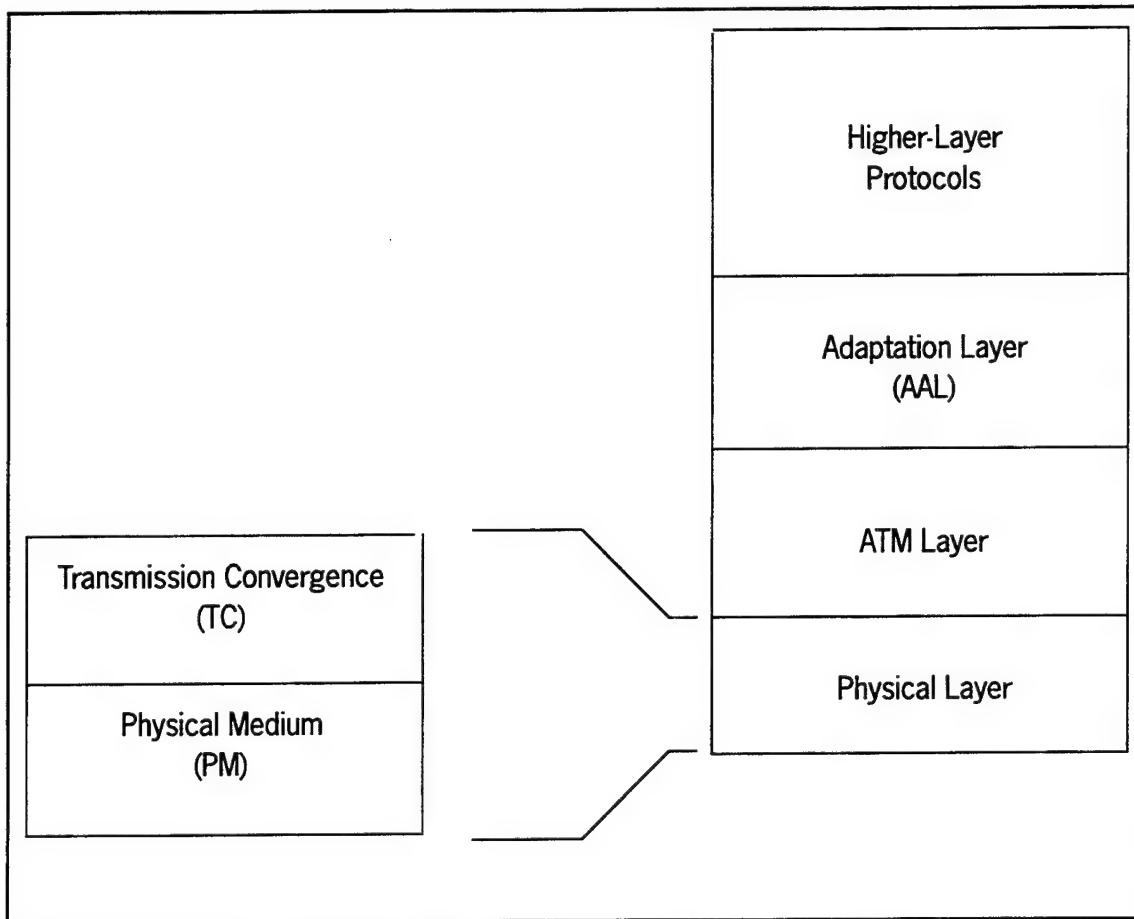


Figure 25. ATM Physical Layer, after [Varma, 1995].

cells are encoded and cell boundaries marked. ATM cells are encapsulated in SONET using an ATM-specific protocol, and standards only exist for SONET speeds of OC-3 and above. [Partridge, 1994]

SONET is often used for framing and synchronization at the physical layer. In addition to the optical media and line rates defined for SONET, the ATM Forum has proposed a variety of physical layer standards such as ATM over twisted-pair wire. [Siu, 1994]

- SONET connections at 155.52 Mbps (OC-3, STS-3)
- DS3 connections at 44.736 Mbps
- 100 Mbps connections using 4B/5B encoding
- 155 Mbps connections using 8B/10B encoding

Figure 26. Physical Interfaces Allowed by the ATM Forum, after [Feibel, 1995].

The Fore NICs used by workstations connected to the NPS ATM LAN are designed for a 62.5/125 μ m multimode fiber with ST connectors running OC-3/SONET [Fore, 1995]. The Cisco A-100 switches in use can support that interface as well as other interfaces listed in Table 3.

Physical Layer	Data Rate(Mbps)	Media	Connector
STS3c/STM1	155.52	Multimode fiber	SC
STS3c/STM1	155.52	Singlemode fiber	SC
STS3c/STM1	155.52	UTP-5	RJ-45
TAXI 4B/5B	100.00	Multimode fiber	MIC
DS3/T3	44.736	Coaxial cable	BNC
E3	34.00	Coaxial cable	BNC

Table 3. Cisco A-100 Interface Types, from [Cisco, 1995].

2. The ATM Layer

The ATM layer provides an interface between the AAL and the physical layer.

The ATM layer is responsible for relaying cells from the AAL to the physical layer, for transmission and from the physical layer to the AAL for use at the end systems. Inside an end system, the ATM layer receives a stream of cells from the physical layer and transmits either cells with new data or empty cells (if there is no data to send). Inside a switch, the ATM layer determines where the incoming cells should be forwarded, resets the corresponding connection identifiers, and forwards the cells to the next link. It also buffers incoming and outgoing cells and handles various traffic management functions such as cell loss priority marking, congestion indication, and generic flow control access. The ATM layer also performs traffic policing by monitoring the transmission rate and conformance to the service contract. [Siu, 1994]

a. ATM Cell Format

The ATM cells are 53 octets long, consisting of a five-octet header and a 48-octet data, or payload, section. ATM cells are not byte oriented. Even though cells are defined as a specific number of octets, the fields within such a cell often cross byte boundaries [Feibel, 1995]. This refusal to align fields with byte boundaries is not efficient from a software designer's viewpoint, essentially ensuring that efficient ATM switch implementations will only be possible using special-purpose hardware. As will be discussed later, the cell header is used by UNIs, NNIs, and switches to route the cell.

Figure 27 shows the ATM cell format.




Figure 27. ATM Cell Format, from [Partridge, 1994].

The payload is formatted in one of the ATM adaptation layer formats described later. The primary reason for CCITT's choice of such a small cell size is to support voice. This came about as a compromise between what the computer industry and the telephone industry each wanted for a cell size. The telephone industry desired a 16-byte cell to deal with the cell-size tradeoffs listed in Table 4. The computer industry desired a 128-byte cell size to minimize the overhead and increase transmission efficiency. The data industry compromised to a 64-byte cell; the telephone industry compromised to a 32-byte cell. Then they split the difference to a 53-byte cell, which is poor for both voice and data. [Varma, 1995]

ATM overhead is large and consumes a significant portion of the bandwidth. A calculation of the overhead in Equation 1 shows that the minimum possible value is 9.4%. This is an extremely large value compared to IP routing overhead for data streams.

- End-to-end delay
- Jitter (delay variation)
- Transmission efficiency
- Switch complexity
- Switching speed
- Buffering requirements
- Echo cancellation requirements
- Quality of conversation requirements

Figure 28. Cell-size Tradeoffs, after [Varma, 1995].

$$\frac{5 \text{ octet header}}{53 \text{ octet header and data}} \times 100\% = 9.4\% \text{ Overhead} \quad (1)$$

ATM adaptation layer (AAL) protocols are discussed later, but it is worth mentioning here that one of the AALs (AAL3/4) consumes four octets from the payload for additional overhead, shown in Equation 2.

$$\frac{5 \text{ octet header} + 4 \text{ octet payload}}{53 \text{ octet header and data}} \times 100\% = 17.0\% \text{ Overhead} \quad (2)$$

b. Interfaces

The fields in the ATM header define the functionality of the ATM layer.

The format of the header for ATM cells has two different forms, one for the user-to-network interface (UNI), and the other for the network-to-node interface (NNI).

[Siu, 1994]

Even though ATM networks are intended to be interconnected to each other, the method of interconnecting depends upon where in the switch hierarchy the connection point is made. UNI connects ATM end-systems (hosts, routers, etc.) to an ATM switch. [Partridge, 1994] identifies UNI's two-fold purpose in Figure 29.

The connections between provider networks are made through NNI. NNI provides smooth interconnection between independently operated ATM networks that trust each other to be well behaved [Partridge, 1994]. More precisely, NNI is any physical or logical link across which two ATM switches exchange NNI protocol. For this reason, the connection between a private ATM switch and a public ATM switch is a UNI, known as a Public UNI, since these switches do not typically exchange NNI information. Because two different connection setups are required, there are two different ATM cell

- To provide an interface to the user equipment that supports multiplexing
- To protect an ATM provider's network from misbehaved user equipment

Figure 29. UNI's Two-fold Purpose, after [Partridge, 1994].

header formats: one for cells crossing an NNI boundary and another for cells given to a UNI by an ATM workstation.

(1) NNI. NNI ATM header is shown in Figure 30. The first two fields in NNI ATM header are the 12-bit virtual path identifier (VPI) and the 16-bit virtual channel identifier (VCI). Both of these are discussed in the following section.

The third field is the 3-bit payload-type field as shown in Figure 31. This field distinguishes between a user cell (when the first bit is zero) and various forms of operations, administration, and management (OAM) cells (when the first bit is set) [Partridge, 1994]. This field allows control and signaling data to be transmitted on a different subchannel from user data, providing separation of user payloads and control data [Siu, 1994].

The payload-type field is a revisiting of the frame relay (FR) standard [Partridge, 1994]. In FR the congestion bit is known as “discard eligible,”

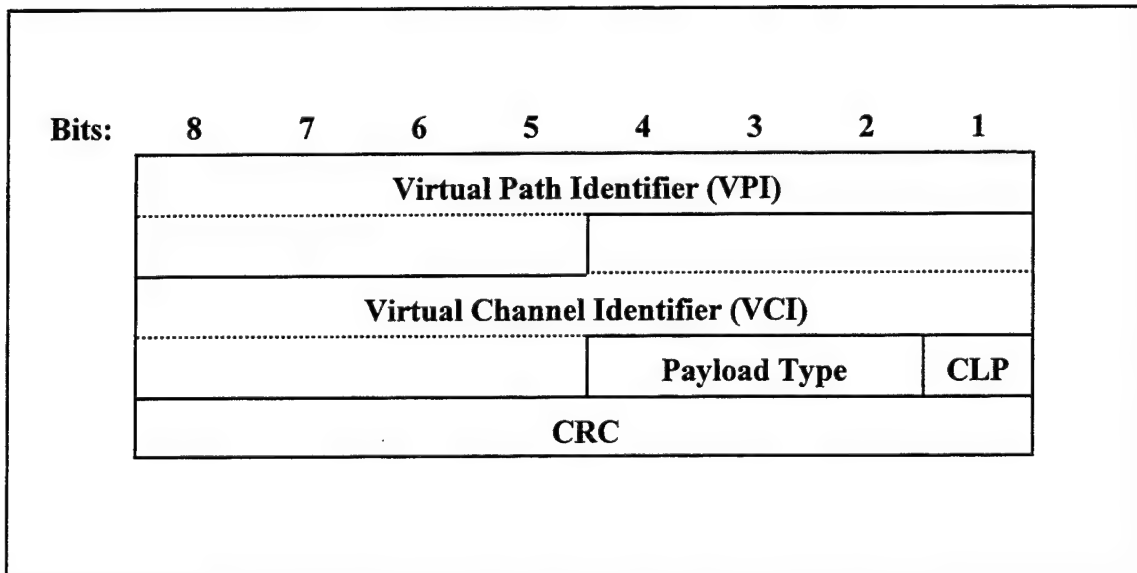


Figure 30. ATM Header at NNI, from [Partridge, 1994].

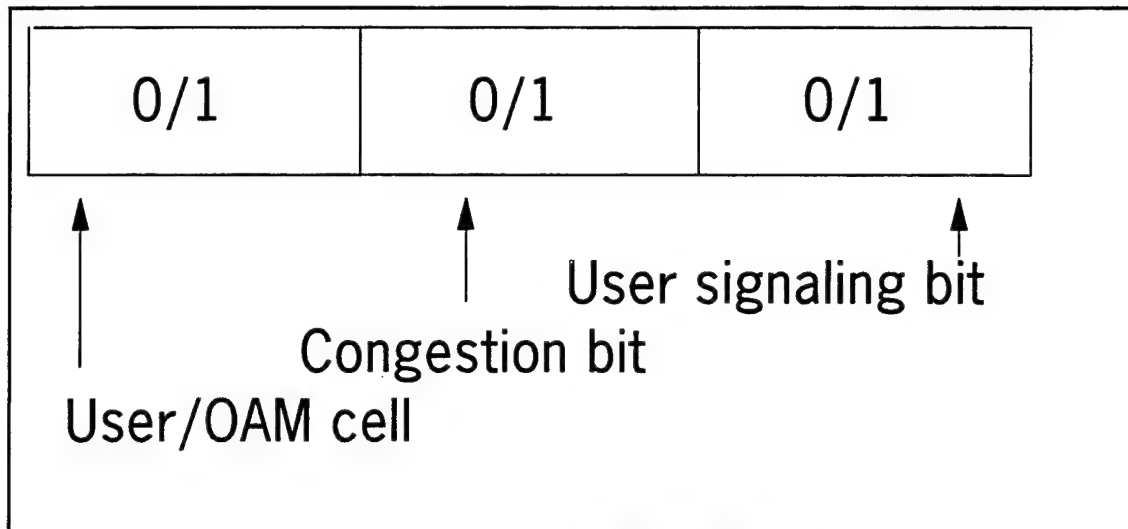


Figure 31. Payload-type Field, from [Varma, 1995].

allowing a congested FR switch to throw a frame away if this bit is set. This field is controlled by the Fore network interface card. Table 4 shows what the bit patterns mean.

Bit Pattern	Definition
000	User data, type 0, no congestion
001	User data, type 1, no congestion
010	User data, type 0, congestion
011	User data, type 1, congestion
100	OAM link associated cell
101	OAM end-to-end associated cell
110	Resource management cell
111	Reserved for future use

Table 4. Payload Type Values, after [Partridge, 1994; Varma, 1995].

The fourth field is the 1-bit cell loss priority (CLP) field. The CLP bit provides the switches with a selective discard capability. This bit could be set by a user to indicate lower-priority cells that could be discarded by the switch during periods of congestion so that the switches on the network can maintain QoS and bandwidth requirements. For example, data applications generally cannot suffer any cell loss without the need for retransmission, but voice and video traffic usually can tolerate minor cell loss. One might therefore set the CLP bit for the noncritical cells in voice or video traffic. The CLP bit can also be used by the network to indicate cells that exceed the negotiated rate limit of a user [Varma, 1995]. This field is controlled by the application software and the network interface card.

The fifth field is the header error check (HEC) field. The HEC field is used to reduce errors in the header that can cause a misrouting of cells from one user into another user's data stream. This field contains the result of an 8-bit CRC on the first four bytes of the 5-byte cell header but not on the data. The CRC is an eighth order polynomial (x^8+x^2+x+1) that can correct a single bit error and detect a large class of multiple bit errors [Freeman, 1991]. Since this header may change at each hop, the CRC will be checked and recomputed at each hop in the ATM network. This overhead is not significant since it is performed by the hardware and virtually eliminates misrouting.

When a switch or an end system terminates the header, multiple-bit errors can be detected and single-bit errors can be corrected. Single-bit errors predominate in all media, but since ATM is intended for use on fiber optic links where

the bit-error rate is less than 10^{-9} , single-bit error correction is quite effective in removing most header errors [Siu, 1994]. Forward error correction requires some overhead in all cases. The Fore NIC and the Cisco A-100 provide HEC in their on-board hardware. [Cisco, 1995; Fore, 1995]

(2) UNI. The ATM header at UNI is slightly different from NNI.

Figure 32 shows UNI header.

The major difference is that the header dedicates four bits to a function called generic flow control (GFC). The purpose of the GFC field is to allow UNI to negotiate with shared access networks about how to multiplex the shared network among the cells of the various ATM connections. Presently these values are not defined and are always zero [Partridge, 1994; Varma, 1995]. It is curious that the single difference between NNI and UNI header formats is unused.

The Fore NIC provides 100 Mbps and 140 Mbps transparent

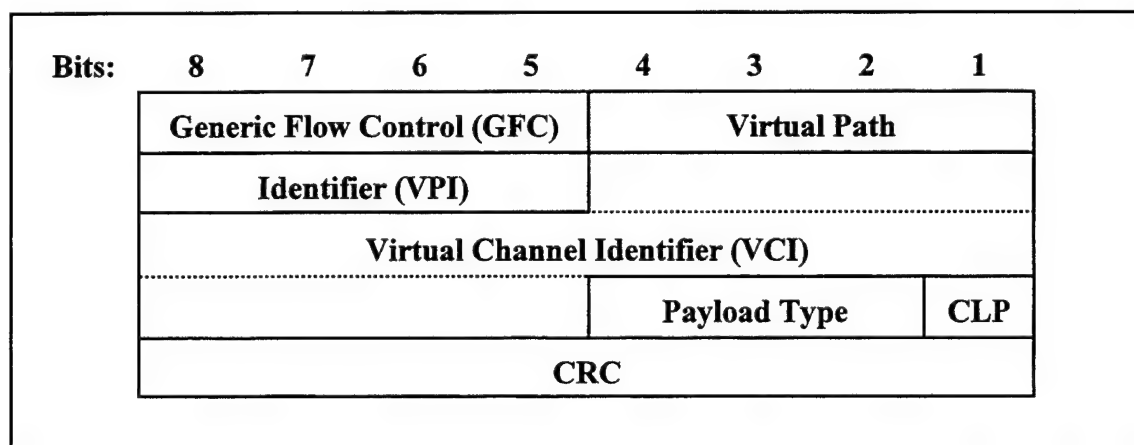


Figure 32. ATM Header at UNI, from [Partridge, 1994].

asynchronous transmitter/receiver interface (TAXI) 4B/5B encoding for OC-3 and SONET on its Intel i960 RISC processor [Fore, 1995]. The Cisco A-100 utilizes an internal 32-bit RISC processor to provide this same functionality [Cisco, 1995]. As TAXI was used mostly for development, the FDDI chipset represents the great majority of the physical layer implementations in use.

ATM network service providers offer several types of interfaces to their networks. One interface is a frame-based interface where one or more of the IEEE 802.x or FDDI frames are supported at UNI, with frame-to-ATM-cell conversion and reassembly being done inside UNI at the source and destination end points, respectively. A gateway host on a LAN might directly connect its Ethernet, token ring, FDDI, or other LAN interfaces to UNI — bridging two widely separated LANs with an ATM backbone network. This preserves the existing investment in these standards and equipment, and enables a gradual transition to ATM networks. [Siu, 1994]

An alternate interface likely to be more popular in the long run, and for which the concept of B-ISDN makes sense, is direct interface at UNI with standard ATM cells. Such a streaming interface can hook subscriber telecom, datacom, and computer equipment directly to the network and may allow orders-of-magnitude greater performance and bandwidth utilization for integrated multimedia traffic of the future. It is for this reason that the IEEE 802.6 packet for the MAC layer of the MAN DQDB protocol looks like an ATM cell. [Siu, 1994]

c. Cell Routing

Because connection setup is required, ATM channel identifiers can be kept short (28 bits). Because ATM is hierarchical, the channel identifiers also have a hierarchical structure. [Partridge, 1994]

Every channel identifier has two parts, as shown in Figure 33.

At the edge of the network, where many transmission links branch out to users' workstations, cells are routed according to the full identifier (VPI and VCI) [Partridge, 1994]. The VPI can be used alone in backbone switches where there are only a few links since the major routing task is determining which link needs to forward the cell [Varma, 1995].

The VPI and VCI together form the routing field which associates each cell with a particular channel or circuit. Where VCI is a single-channel identifier, a virtual path is a bundle of virtual channels, all of which are switched transparently across the ATM network on the basis of the common VPI [Alles, 1995]. However, the VPIs and VCIs have local significance only on a particular link; the contents of the routing field will generally change as the cell traverses from link to link, being remapped as

- Virtual path — identified by the virtual path identifier (VPI)
- Virtual channel — identified by the virtual channel identifier (VCI)

Figure 33. Channel Identifiers, after [Alles, 1995].

appropriate at each switch [Siu, 1994]. In normal operation, switches allocate all UNI connections within VPI=0 [Alles, 1995], as does the NPS ATM LAN [Cisco, 1995].

The 28-bit VPI/VCI combination uniquely identifies the ATM connection and provides a two-level addressing and routing hierarchy. The ATM Forum made the decision to go with a small address space in order to make the switch-to-switch routing easier — VPIs/VCIs are small enough to be used as indices into routing tables. Since the address space is small — for example, the ISO network address space can be up to 160 bits long — the VPIs/VCIs are required to be unique on a hop-by-hop basis between switches. [Partridge, 1994]

For UNI the routing field contains 24 bits, and at NNI the routing field contains 28 bits. This means that the number of sessions that can be active at a UNI has been limited to slightly greater than 16 million (2^{24}) compared to slightly greater than 268 million (2^{28}) for NNI. The Cisco A-100 switch supports 4096 (2^{12}) channels per line [Cisco, 1995].

The Cisco A-100 switch's line interface (LINF) card receives a cell sent over a medium. A header translator performs the HEC and generates internal switch-specific overhead (SSO) information using VPI, VCI, and payload type (PT) in the cell header. If the cell belongs to a point-to-point connection, the LINF inserts new VPI and VCI values in the cell header, as discussed in the next section. The SSO transfers to the expandable ATM output-buffer modular switch (XATOMSW) along with the cell. The XATOMSW switches and sends the cell and SSO to the destination LINF according to

the SSO information. When a specific line is congested, the system temporarily saves overflow cells in a 2048-cell input buffer and a 128-cell output buffer. Two lines share a buffer pool [Cisco, 1995]. A full discussion of switch construction and buffering considerations appears in [Varma, 1995].

The LINF inserts new VPI and VCI values in the cell header according to information in the SSO after receiving the cell and SSO if the cell belongs to a point-to-multipoint connection. This series of events results in cell transmission to the destination line. [Cisco, 1995]

3. ATM Adaptation Layer (AAL)

The topmost layer of protocols for packaging data into cells is collectively referred to as the ATM adaptation layer (AAL), and the individual protocols are referred to as AALs. The purpose of the ATM adaptation layer is to efficiently package the various kinds of higher level data into a series of cells that can be sent over ATM connections and reconstructed into the appropriate format at the receiving end [Partridge, 1994]. AAL provides the necessary protocol translation between ATM and other communication services (such as voice, video, or data) involved in a transmission. For example, the AAL translates elements from a pulse-code modulation (PCM) transmission (which encodes voice data in digital form) into ATM cells [Feibel, 1995]. CCITT proposed four types of AALs, each supporting a different type of traffic or service expected to be used on ATM networks, according to the properties listed in Figure 34.

- The information being transported is time dependent/independent. It may be necessary to regenerate the time dependency of a signal at the destination, e.g. a 64 Kbps PCM voice.
- Variable/constant bit rate.
- Connection-oriented/connectionless mode information transfer.

Figure 34. Traffic or Service Proposed by CCITT, after [Varma, 1995].

These properties define eight possible classes, four of which are specifically defined as B-ISDN service classes. The four ATM adaption layer services shown in Table 5 are defined to match up with the four B-ISDN information classes, shown in Figure 35.

Due to initial performance, overhead, and complexity of AAL1, 2, 3, and 4, a new layer (AAL5) was subsequently standardized [Partridge, 1994]. AAL5 supports classes C or D more efficiently than AAL3/4 [Feibel, 1995]. The three-fold goals of AAL5, also

Traffic Class	A	B	C	D
Timing between source and destination	Yes		No	
Bit rate	Constant	Variable		
Connection mode	Connection-oriented			Connectionless
AAL	1	2	3/4, 5	3/4, 5

Table 5. AAL Types, from [Varma, 1995].

- *Class A — Constant bit rate (CBR) service:* AAL1 supports a connection-oriented service in which the bit rate is constant. Examples of this service include 64 Kbps voice, fixed-rate uncompressed video, and leased lines for private data networks.
- *Class B — Variable bit rate (VBR) service:* AAL2 supports a connection-oriented service in which the bit rate is variable but requires a bounded delay for delivery. Examples of this service include compressed packetized voice or video.
- *Class C — Connection-oriented data service:* The ITU originally recommended two types of AAL protocols (AAL3 and AAL4) to support this service class, but these two types have been merged into a single type — AAL3/4. Because of the high complexity of AAL3/4 protocols, the AAL5 protocol was proposed and is often used to support this class of service. Examples of this service include connection-oriented file transfer and data network applications where a connection is set up before data is transferred. This service has variable bit rate and does not require bounded delay for delivery.
- *Class D — Connectionless data service:* Either AAL3/4 or AAL5 can be used to support this class of service. Examples of this service include datagram traffic and data network applications where no connection is set up before data is transferred.

Figure 35. Four B-ISDN Service Classes, after [Siu, 1994].

known as the simple and efficient adaptation layer (SEAL), are shown in Figure 36.

Although each AAL design is optimized for a specific type of traffic, there is no stipulation in the standards that AALs designed for one class of traffic cannot be used for another. In fact many vendors of ATM equipment currently manufacture products that exclusively use AAL5 to support all the above classes of traffic. Most activities at the ATM Forum have focused on AAL5.

- Have low overhead
- Minimize the computer's cost in handling cells
- Behave as much as possible like existing data communication interfaces — Ethernet and FDDI

Figure 36. Goals of AAL5 (SEAL), after [Partridge, 1994].

The NPS ATM LAN's Cisco A-100 switches support AAL1 through AAL5 as well as all traffic types. The Fore NICs, however, support AAL3/4 and AAL5 only. Since applications at NPS run AAL5 exclusively, that protocol is discussed below. The details of the other protocols can be found in [Partridge, 1994; Suzuki, 1994; Varma, 1995].

CCITT divided each AAL into two sublayers, shown in Figure 37. A separate

- *SAR (segmentation and reassembly)* — the lower sublayer that packages variable-sized packets into fixed-sized cells at the transmitting end and repackages the cells at the receiving end. The SAR sublayer is also responsible for finding and dealing with cells that are out of order or lost.
- *CS (convergence sublayer)* — the upper sublayer that wraps the user-service data units in a header and trailer which contain information used to provide the services required. The information in the header and trailer depends on the class of information to be transported but will usually contain error handling and data priority preservation information. CS provides the interface for the various services. Users connect to the CS through service access points (SAPs).

Figure 37. AAL Sublayers, after [Feibel, 1995].

protocol data unit (PDU) is defined for each class of service. Each PDU contains 48 octets that are allocated for the header, trailer, and payload. These PDUs become the payload for the ATM cells that are transmitted [Feibel, 1995]. Figure 38 shows the AAL sublayers.

The 1-bit SAR sublayer is encoded in the last bit of the payload-type field of the ATM header. The receiving computer queues cells until it receives a cell with the end-of-packet bit set. Figure 39 shows the AAL5 SAR format.

The AAL5 convergence sublayer fills the last eight bytes of the final cell with four

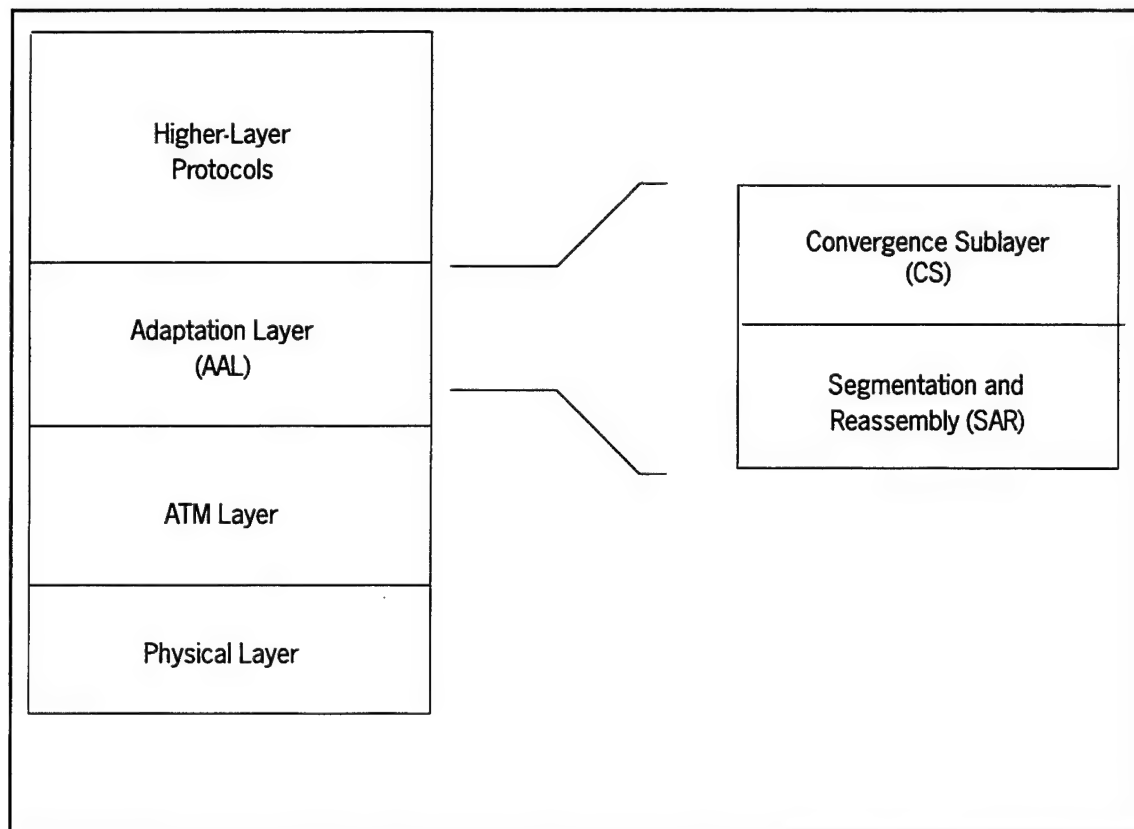


Figure 38. ATM Adaptation Layer (AAL), after [Varma, 1995].




Figure 39. AAL 5 SAR Format, from [Partridge, 1994].

fields that manage the SAR. This final cell has 40 bytes of data and/or padding and has eight bytes of trailer. Figure 39 shows the AAL5 convergence sublayer.

The first field is an 8-bit user-to-user indication (UU) field. The second field is an 8-bit common part indicator (CPI) field. Both fields are currently unused and set to zero. [Partridge, 1994; Varma, 1995]

The third field is a 16-bit length field that states the number of bytes of data in the packet. This field is used for reassembling the original packet from the ATM cells.

The fourth field is an extremely robust 32-bit CRC. This CRC checks the entire convergence sublayer packet, including the pad and trailer, and can detect both lost and misordered cells. [Partridge, 1994]

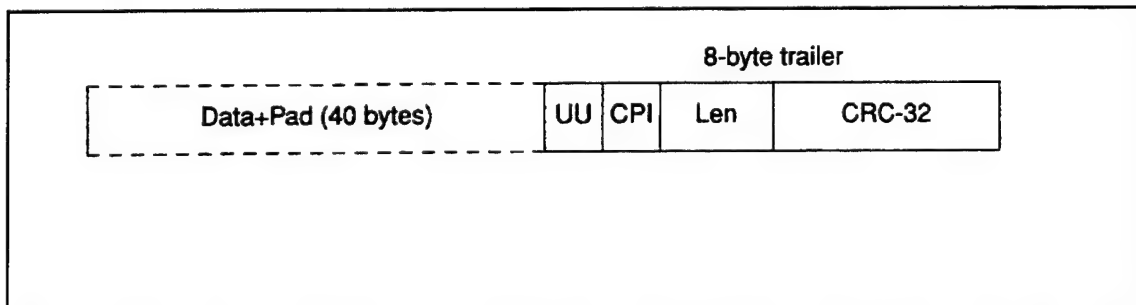


Figure 40. AAL5 Convergence Sublayer Format, from [Partridge, 1994].

The Fore NIC does all of the SAR and CS functions in its dedicated embedded Intel i960 RISC processor. ATM switches do not deal with the AAL level.

D. ATM SIGNALING

The basic operation of an ATM switch is simple: receive a cell across a link on a known VCI/VPI, look up the connection value in a local translation table to determine the outgoing port (or ports) of the connection on that link, then retransmit the cell on that outgoing link with the appropriate connection identifiers. [Alles, 1995]

The switch operation is simple because external mechanisms set up the local translation tables prior to the transmittal of any data. The manner in which these tables are set up determines the two fundamental types of ATM connections, shown in Figure 41.

The Fore NICs and the Cisco A-100 support both PVCs and SVCs. As was noted in Chapter IV, the PacBell switches also support both PVCs and SVCs, but the Sprint

- *Permanent Virtual Connections (PVC)* — a connection setup by network management in which a set of switches between the ATM source and destination system are programmed with the appropriate VPI/VCI values. PVCs always require some manual configuration, therefore their use can be cumbersome.
- *Switched Virtual Connections (SVC)* — a connection that is set up automatically through a signaling protocol. SVCs do not require the manual interaction needed to set up PVCs and are expected to be more widely used. All higher layer protocols operating over ATM primarily use SVCs.

Figure 41. Two Fundamental Types of ATM Connections, after [Alles, 1995].

switches support PVCs only.

Before ATM cells can be transmitted, an ATM connection must be established between the sender and receiver, and a VPI/VCI must be assigned to the connection at each hop. The protocol that performs these tasks is called a signaling or setup protocol. Designing and implementing a signaling protocol is one of the hardest and most complex tasks in networking. [Partridge, 1994]

1. Connection Management

ATM signaling is initiated by an ATM end-system that desires to set up a connection through an ATM network, and signaling packets are sent on a well known virtual channel: VPI=0, VCI=5. The signaling is routed through the network, from switch to switch, setting up the connection identifier as it goes, until it reaches the destination end system. The destination can either accept and confirm the connection request, or else can reject the requesting signal, clearing the connection. Note that because the connection is set up along the path of the signaled connection request, the data also flows along this same path. [Alles, 1995]

ATM signaling uses the “one-pass” method of connection set-up, which is the model used in all common telecommunications networks. A connection request from the source end-system is propagated through the network, setting up the connection as it goes, until it reaches the final destination end system. The routing of the connection request — and hence any subsequent data flow — is governed by the ATM routing protocols. These protocols route the connection request based upon both the destination

address and the traffic and QoS parameters requested by the source. The destination may choose to accept or reject the connection request. Since call routing is based purely on the parameters in the initial connection request message, negotiation of connection parameters is limited. [Alles, 1995]

The Q.93B signaling specification supports the creation of three types of connections. A requester can establish a unidirection channel to a destination, a two-way channel to a destination (but only if the bandwidth in both directions is identical), or a unidirectional multicast delivery tree.

As currently defined, Q.93B is simply a set of rules for how the sending and receiving ends of a virtual circuit negotiate connectivity with the ATM network. How switches within the ATM network communicate with each other is not standardized. ATM switches can use Q.93B or their own protocol. However, whichever protocol is used, Q.93B defines the services the setup protocols must offer. [Alles, 1995]

The exchange of Q.93B messages required to establish a simple one-way connection between a sender and receiver is shown in Figure 42. The sender begins by sending a SETUP message. The SETUP message contains information about the call, including a low specification, which AAL is to be used (and AAL-specific parameters like the largest packet that can be sent if using AAL5), and the address of the receiver. The network acknowledges the SETUP message with a CALL PROCEEDING message. The primary purpose of this message is to acknowledge the SETUP message but it also contains the network-assigned VCI/VPI for the connection.

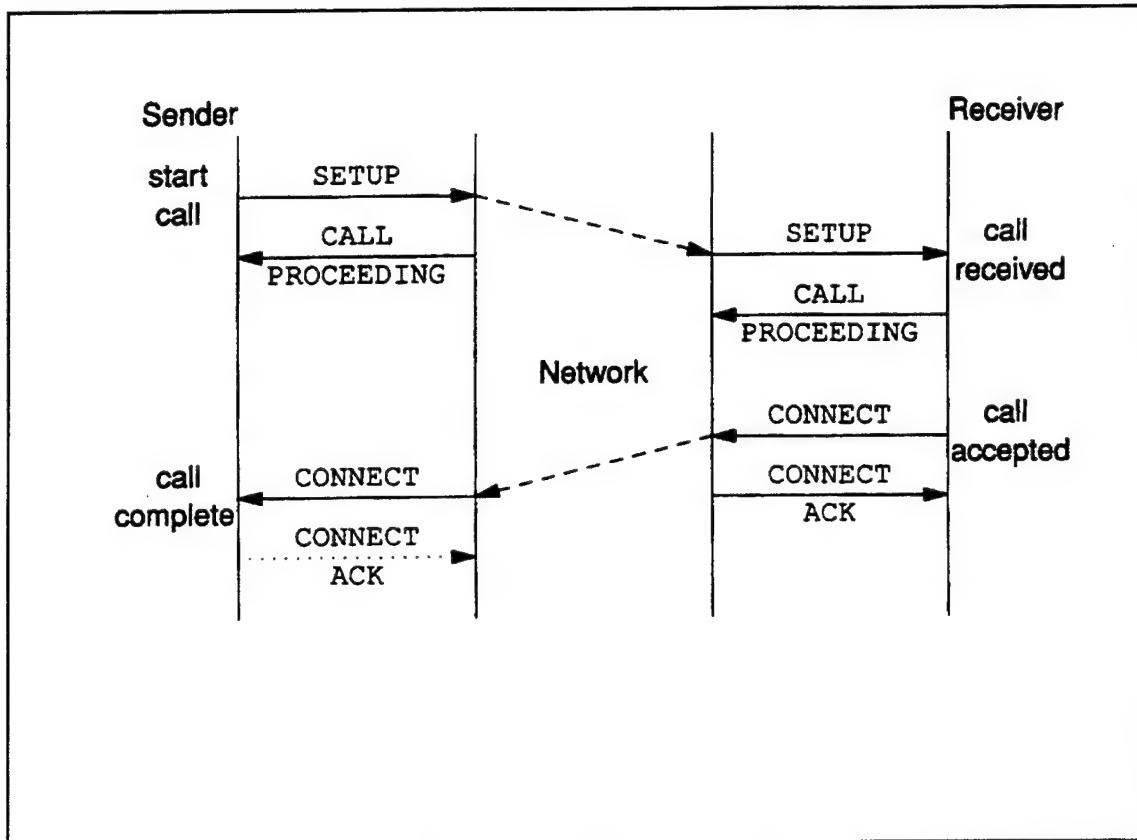


Figure 42. Q.93B Signaling Connection Setup, from [Partridge, 1994].

The network notifies the receiver that a call is being requested by sending a SETUP message to the receiver. The receiver acknowledges with a CALL PROCEEDING message and then decides, based on information in the SETUP message, whether to accept the call. If the receiver decides to accept the call, it sends a CONNECT message to the network. The network acknowledges the receiver's message with a CONNECT ACKNOWLEDGE and in turn sends a CONNECT to the sender indicating the call is established. If there are errors in the negotiation, the typical response is for the party that detects the error to send a RELEASE COMPLETE message which indicates that the call has failed and all information on the call must be discarded.

To ensure reliable transmission of messages between the endpoints and the network, the SETUP and CONNECT messages are retransmitted if they are not acknowledged after a certain reasonable time interval. It is not permissible to send data until a CONNECT message is received, so if the network detects the sender transmitting data on the new VCI, this information is taken to be an implicit acknowledgment that the CONNECT message has been received. [Alles, 1995]

For example, assume we have an ATM network as shown in Figure 43 with ATM switches at nodes A, B, C, D, E, and F. Assume that we have ATM UNI cell interfaces at hosts AA and FF.

This is what happens: host AA makes a connection request to UNI at switch A. After an exchange of connection parameters between AA and UNI (such as destination, traffic type, peak and average bandwidth requirement, delay and cell loss requirement,

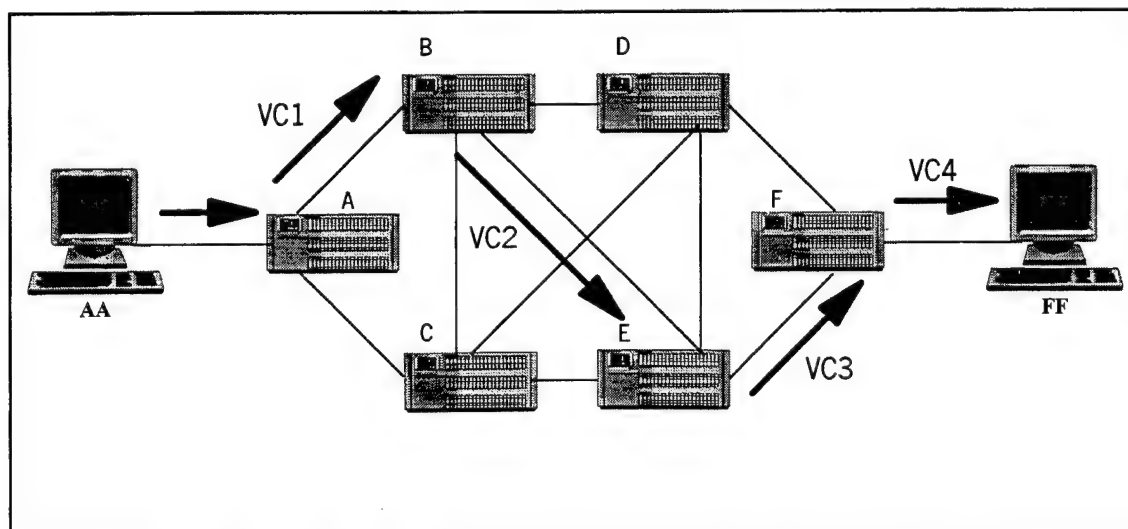


Figure 43. Sample ATM Switching Network, after [Ebrahim, 1992].

cost function, etc.), UNI at switch A forwards the request to the network. The software running on the network computes a route based on the cost function specified by host AA, and figures out which links on each leg of the route can best support the requested quality of service and bandwidth. Then switch A sends a connection setup request to all the nodes in the path en route to the destination node in host FF.

Let us assume that the ATM switch route selected was A--B--E--F. Each of these four nodes picks an unused VCI label and reserves it for the connection in its connection lookup table inside. Say that switch A picks VC1. It will forward VC1 to switch B. Switch B in turn picks output VC2, associates it with input VC1 in its connection table, and forwards VC2 to switch E (see example in Table 6 for switch B). Switch E picks output VC3 and associates it with input VC2 in its connection table and forwards VC3 to switch F. Switch F picks output VC4, associates it with input VC3 in its connection table, and queries the addressed UNI to see if it would accept this connection request.

If host FF returns an affirmative to this query, switch F hands UNI (at switch F) and host FF the value VC4 as a connection identifier for this connection. Switch F then sends an acknowledgment (ACK) back to switch E. Switch E ACKs back to switch B and sends it VC3. Switch B puts VC3 in its connection tables to identify the path going in the reverse direction, and ACKs to switch A sending it VC2. Switch A associates VC2 in its connection tables with VC1, and ACKs the originating UNI with VC1. UNI hands host AA VC1 and the connection is finally established [Ebrahim, 1992]. Two table entries have been created in each switch for traffic going in each direction.

Incoming		Outgoing	
Port	VC Number	Port	VC Number
A	1	C	4
	3	D	6
	5	E	2
C	4	A	3
	7	D	8
	9	E	10
D	6	A	5
	8	C	7
	11	E	12
E	2	A	1
	10	C	9
	12	D	11

Table 6. Switch B's Routing Table. Boldface values show table entries for example shown in Table 5.

Host AA identifies the connection with VCI label VC1, and host FF identifies the connection with VCI label VC4. In this extended example, the labels get translated at each node to the next outgoing label as shown in Table 7.

As mentioned in Figure 43, each switch maintains a routing table of ports and VCI/VPIs. An example of a routing table for switch B in the previous example is shown in Table 7.

Host		Switches				Host	
A		B	E	F			
AA	⇒ VC1(UNI)	⇒ VC2	⇒ VC3	⇒ VC4(UNI)	⇒	FF	
AA	⇐ VC1(UNI)	⇐ VC2	⇐ VC3	⇐ VC4(UNI)	⇐	FF	

Table 7. Example VC Table Mappings, after [Ebrahim, 1992].

Other scenarios are also possible and can depend on a vendor's implementation of the ATM network. When host AA or host FF terminate the connection, it is "torn down" via signaling messages, and the VCI labels can be reused for other connections [Ebrahim, 1992]. The NPS ATM LAN provides signaling, guaranteed bandwidth reservation, and cell loss priority by the Fore NICs and the Cisco A-100 switches. [Cisco, 1995]

What conclusions may a user attachment make, given a VCI label? As is shown in the above example, none. The VCI labels are owned by network nodes and get randomized quite quickly as connections come and go. It is possible to have certain reserved VCI labels as identifiers for special, well-known services that may be provided by the network, but little can be assumed about the dynamically assigned VCI labels for most user related connections. [Ebrahim, 1992]

2. Connection Types

One major key to understanding ATM is to realize that there are two fundamental ATM connection types, shown in Figure 44.

Missing from these two types of ATM connections is an analog to the multicasting or broadcasting capability common in many shared medium LAN technologies. In these dissimilar technologies, multicasting allows multiple end systems to both receive data from other multiple systems and to transmit data to these multiple systems. Such capabilities are easy to implement in shared media technologies such as LANs where all nodes on a single LAN segment must necessarily process all packets sent on that segment. The obvious analog in ATM to a multicast LAN group is a bidirectional, multipoint-to-multipoint connection. Unfortunately this obvious solution cannot be implemented when using AAL5, and in any case does not scale well for connection-oriented circuits when many participants are senders.

Unlike AAL3/4 with its Message Identifier (MID) field [Forum, 1994], AAL5

- *Point-to-point connections* — connect two ATM end-systems. Such connections can be unidirectional but are typically bidirectional.
- *Point-to-multipoint connections* — connect a single source end-system (the root node) to multiple destination end-systems (leaves). Cell replication is done within the network by the ATM switches at which the connection splits into two or more branches. Such connections are unidirectional, permitting the root to transmit to the leaves, but not the leaves to transmit to the root, or to each other, on the same connection.

Figure 44. ATM Connection Types, after [Alles, 1995].

does not have any provision within its cell format for the interleaving of cells from different AAL5 packets on a single connection. This means that all AAL5 packets sent to a particular destination across a particular connection must be received in sequence, with no interleaving between the cells of different packets on the same connection, otherwise the destination reassembly process is unable to reconstruct the packets.

This restriction is the reason ATM AAL5 point-to-multipoint connections can only be unidirectional, for if a leaf node is to transmit an AAL5 packet onto the connection, it can be received by both the root node and all other leaf nodes. However, at these nodes, the packet sent by the leaf can well be interleaved with packets sent by the root and possibly other leaf nodes. This out-of-order reception can preclude the reassembly of any of the interleaved packets. [Alles, 1995]

Clearly this lack of a multipoint-to-multipoint multicast capability in ATM is not acceptable since most existing protocols rely upon the existence of a low-level multicast/broadcast facility. Figure 45 shows three methods proposed for solving this problem.

The overlaid mechanism requires each node to maintain N connections for each group, where N is the total number of transmitting nodes within the group, while the multicast server mechanism requires only two connections. The overlaid mechanism also requires a registration process for telling nodes that join a group what the other nodes in the group are, so that it can form its own point-to-multipoint connection. The other nodes

- *VP-multicasting* — a multipoint-to-multipoint VP links all nodes in the multicast group, and each node is given a unique VCI value within the VP. Interleaved packets can be identified by unique VCI value of the source. Unfortunately, this mechanism requires a protocol to uniquely allocate VCI values to nodes. Such a mechanism does not currently exist, and current SAR devices do not support such a mode of operation.
- *Multicast server* — all nodes wishing to transmit onto a multicast group set up a point-to-point connection with an external device known as a multicast server (a resequencer or serializer). The multicast server is connected to all nodes wishing to receive the multicast packets through a point-to-multipoint connection. The multicast server then retransmits them across the point-to-multipoint connection after ensuring that the packets are serialized (that is, one packet is fully transmitted prior to the next being sent). This precludes cell interleaving. This also adds a bottleneck prone to single-point failure, and imposes performance limitations inhibiting fast leave/join, and limiting the maximum number of participants.
- *Overlaid point-to-multipoint connections* — all nodes in the multicast group establish a point-to-multipoint connection with each other node in the group, and becomes a leaf in the equivalent connections of all other nodes. All nodes can both transmit to and receive from all other nodes. Creating such fully connected meshes requires n^2 connections for n participants in each channel. Previously noted server constraints also pertain since a server is required to keep track of all participants.

Figure 45. Proposed ATM Multicasting Methods, after [Alles, 1995].

also need to know about the new node so they can each add the new node to their own individual point-to-multipoint connections.

The multicast server mechanism is more scalable in terms of connection resources, but has the problem of requiring a centralized sequencer. The resequencer is both a potential bottleneck and a potential single point of failure.

As yet there is no practical solution for many-to-many ATM multicast. Higher layer protocols within ATM networks use both of the latter two solutions for multicast. This is one example of why adapting existing internetworking protocols to ATM is so complex. Most current protocols, particularly those developed for LANs, implicitly assume a network infrastructure that is a shared-medium connectionless technology with implicit broadcast mechanisms. As noted above, ATM violates all of these common assumptions. [Alles, 1995]

The Fore NICs and the Cisco A-100 switches both support multicasting with dynamic addition and deletion of recipients without throughput degradation using the ATM Forum's overlay model. The Cisco A-100 can support up to 4096 point-to-point connections per interface and 1024 point-to-multipoint connections per switch [Cisco, 1995; Fore, 1995]. Thus a total of $(1024)^{1/2} \approx 32$ participants might be accommodated in a single multicast channel ATM exercise. We have not yet tested this capability. We are currently able to handle 500 - 1000 simultaneous participants in global IP multicast exercises, with the constraining bottleneck being the T1 (1.5 Mbps) shared campus connection. Clearly these constraints on multicasting over ATM are significant. They are not well understood or widely acknowledged in the ATM community.

3. Uni 3.0

ATM signaling protocols vary by the type of ATM link. ATM UNI signaling is used between an ATM end-system and an ATM switch across an ATM UNI, and ATM NNI signaling is used across NNI links. Standards partly exist for ATM UNI signaling,

although work is continuing on NNI signaling [Alles, 1995]. A standard for ATM UNI signaling is described in the ATM Forum UNI 3.1 specification [Forum, 1994], which is a slight modification to the earlier UNI 3.0 specification [Forum, 1993]. UNI signaling requests are carried across UNI in a well-known default connection: VPI=0, VCI=5.

Apart from some minor "bug" fixes, the only substantive difference between UNI 3.0 and UNI 3.1 appears in the data link protocol called the service specific convergence protocol (SSCOP). SSCOP is used for the reliable transport of the ATM signaling packets. UNI 3.1 brought the ATM Forum signaling specification into alignment with the International Telecommunications Union and Telecommunications (ITU-T) sector's Q.2931 signaling protocol stack. UNI 3.0 had referenced the earlier draft, Q.93b. Although there are no functional differences between UNI 3.0 and UNI 3.1, the two are not interoperable due to the differences in the data link protocol. UNI 3.0 referenced an earlier, non-interoperable draft of Q.2100 — Q.SAAL.

UNI 3.1 specification is based upon Q.2931, a public-network signaling protocol developed by ITU-T that was based upon the Q.931 signaling protocol used with narrowband ISDN (N-ISDN). The ATM signaling protocols run on top of SSCOP, which is a data link protocol that guarantees delivery through the use of sliding time windows and retransmissions. [Alles, 1995]

It is important to note that ATM per se does not offer an assured service. Cells are not retransmitted by ATM devices upon loss since higher layers (such as TCP) will handle reliable delivery. This makes ATM devices simpler, faster, and cheaper. ATM

signaling, however, does require assured delivery guarantees from SSCOP since it does not run on any standard higher layer protocol like TCP, and the intermediate signaling state machines can be made much simpler if assured delivery is assumed [Alles, 1995]. Refer to [Partridge, 1994] for further discussion on reliable delivery in ATM networks.

The greatest contribution of UNI 3.0/3.1 was the addressing structure. Any signaling protocol requires an addressing scheme that allows the signaling protocol to identify the source and destination of connections. The ITU-T had long settled upon the use of the E.164 telephone-number-like addresses as the addressing structure for public ATM (B-ISDN) networks. This is equivalent to Ethernet addressing, not IP addressing. Since E.164 addresses are a public resource (and cannot typically be used within private networks) the ATM Forum extended ATM addressing to include private networks. [Alles, 1995]

More specifically, in developing a private network addressing scheme for UNI 3.0/3.1, the ATM Forum settled on the *subnetwork* or *overlay* model. This model decouples the ATM layer from any other existing protocol, defining an entirely new addressing structure. All existing protocols are expected to operate over the ATM network. Figure 46 shows the overlay model's ATM addressing scheme.

The overlay model requires the definition of both a new addressing structure and an associated routing protocol. All ATM systems will then need to be assigned an ATM address in addition to any higher layer protocol addresses it will also support. The ATM addressing space can be logically disjoint from the addressing space of whatever protocol

is running over the ATM layer and typically does not bear any relationship to it. All protocols operating over an ATM subnet also require some form of ATM address resolution protocol (ATM ARP) to map higher layer addresses (such as IP addresses) to their corresponding ATM addresses [Alles, 1995]. The Fore NICs and the Cisco A-100 switches implement Uni 3.0 signaling (Q.2931) only. [Cisco, 1995; Fore, 1995]

4. NSAP

The ATM Forum defined an address format for private networks based on OSI network service access point (NSAP). It is important to note that an ATM address is not an NSAP, despite the similar structure, because they do not identify NSAPs but rather subnetwork points of attachment. While such addresses are commonly referred to as "NSAP addresses," [Alles, 1995] describes them better as ATM private network addresses or ATM end-point identifiers.

The 20-byte NSAP is designed for use within private ATM networks, while public networks typically use E.164 addresses that are formatted as defined by ITU-T. The ATM Forum did specify an NSAP encoding for E.164 addresses. This will be used for encoding E.164 addresses within private networks but may also be used by some private networks. Such networks may base their own NSAP format on the E.164 address of the public UNI to which they are connected and take the address prefix from the E.164 number, identifying local nodes by the lower order bits.

All NSAP format ATM addresses consist of three components: an authority and format identifier (AFI), which identifies the type and format of the initial domain

identifier (IDI); the IDI, which identifies the address allocation and administration authority; and the domain specific part (DSP), which contains actual routing information. The Q.2931 protocol defines source and destination address fields for signaling requests, and also defines subaddress fields for each. [Alles, 1995]

Three formats of private ATM addressing are discussed in Figure 46. They differ by the nature of the AFI and IDI.

The ATM Forum recommends that organizations or private network service providers use either the DCC or ICD formats to form their own numbering plan. Organizations that want to obtain ATM addresses can do so through the same mechanism used to obtain NSAP addresses: ANSI [Alles, 1995]. Once obtained, such addresses can be used for both ATM addresses and for NSAP addressing. The Fore NICs implement the ICD ATM format [Fore, 1995]. Figure 47 shows the ICD ATM format.

The DSP is typically subdivided into a fixed hierarchy that consists of a routing

- NSAP Encoded E.164 format — the IDI is an E.164 number.
- DCC Format — the IDI is a data country code (DCC); these identify particular countries, as specified in ISO 3166. Such addresses are administered by the ISO National Member Body in each country.
- ICD Format — the IDI is an international code designator (ICD); these are allocated by the ISO 6523 registration authority (the British Standards Institute). ICD codes identify particular international organizations.

Figure 46. Three Formats of Private ATM Addressing, after [Alles, 1995].

domain (RD), an area identifier (AREA), and an end system identifier (ESI). The ATM Forum, however, has combined the RD and AREA fields into a single high-order DSP (HO-DSP) field which is then used to support flexible, multilevel addressing hierarchies for prefix-based routing protocols. No rigid boundary exists within the HO-DSP. Instead a range of addressing hierarchies will be supported using prefix masks [Alles, 1995]. A discussion of how the NSAP address for the NPS ATM LAN were established is provided in Appendix D. A full discussion of the coexistence of IPv6 and ATM can be found at [Jarrin, 1996].




Figure 47. ICD ATM Format, from [Alles, 1995].

E. SUMMARY

This chapter outlines the major software considerations pertinent to the NPS ATM LAN. It reviews the theoretical underpinnings of ATM, and how the Fore NICs and the Cisco A-100 switches meet the requirements of ATM. Then it discussed the general considerations in building the LAN. Appendices D and E fully outline the details for formatting and configuring the Fore Cards and the Cisco switches.

VI. EXPERIMENTAL RESULTS

A. INTRODUCTION

This chapter provides the results of experimental testing performed on the NPS ATM LAN. The testing is done by using standard Unix commands. The tests are performed for each of the five configurations during the LAN's evolution: stand-alone configuration, then peer-to-peer, through one switch, through two switches, and finally across campus and out to BayNet. Topologies for each configuration are detailed in Chapter IV. The overall failure to establish a sustainable BayNet ATM WAN is also examined.

B. SOFTWARE TEST DESCRIPTION

A program created at the Army's ballistic research lab (BRL) called *nttcp* is used for most benchmark testing [BRL, 1984]. *nttcp* is an analysis tool that runs both TCP and UDP streams for testing connectivity and performance between ATM or Ethernet interfaces on two systems. It differs from common "blast" tests, which usually do not allow measurements at the remote end of a UDP transmission. *nttcp* creates a binary large object (BLOB) file of 134,217,728 bytes (134.2 Mbytes) and uses ftp to transfer it between the designated workstations. Once the ftp is completed, the file is destroyed. The file is large enough to be useful in determining the average transfer rate.

This program is run in both receive and send modes. One host acts as sender and the other as receiver. On the receiving host from the Unix command prompt run *nttcp -r*;

in the sending host run *nttcp -t destination_hostname*. This default setting is to transfer a TCP stream, but by use of the *-u* switch, a UDP stream can be transferred.

C. CONFIGURATIONS

1. Stand Alone

As described in Chapter IV, the first step in the installation process is to select two Silicon Graphics workstations in STL on which to build the initial portion of the ATM LAN. Figure 48 shows these two workstations in the standalone configuration. A Fore NIC is installed and configured in each of these workstations as described in Appendix D.

nttcp benchmark testing is performed from Royal to Royal. The following Unix line commands are issued to a Royal command window:

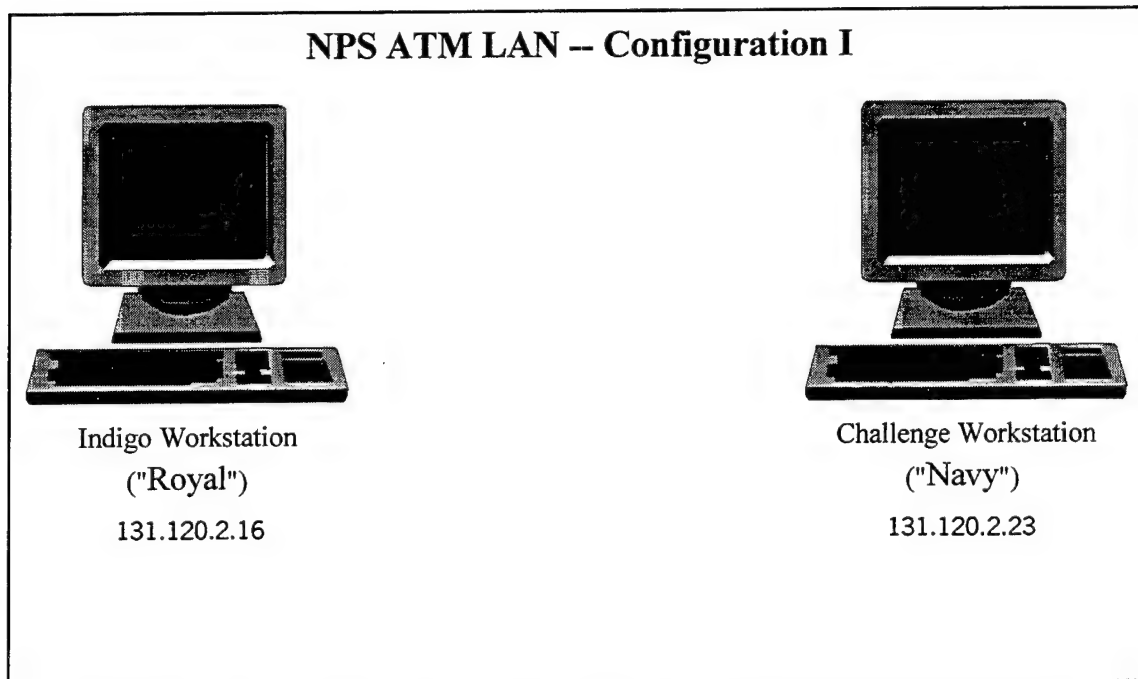


Figure 48. Stand Alone.

nttcp -r

nttcp -t royal

with a result of 132 Mbps with no load on Royal. Backplane signaling speed and CPU load are the principal constraints on transfer rate.

nttcp benchmark testing is also performed from Navy to Navy. The following commands are issued to a Navy command window:

nttcp -r

nttcp -t navy

with a result of 89 Mbps with a large cpu load on Navy. Navy is the file server for STL, so a slower rate and a larger cpu load is normal.

The surprising discovery from these tests is that the workstations are I/O bound with respect to the ATM traffic. Even when there is neither a switch nor a direct ATM connection to another workstation, the CPU load on the workstations will drastically control the data rate and throughput. The Fore helpdesk verified this situation to be the case. Even though the switches can support up to 155 Mbps, the actual throughput will be much less and will depend upon end host CPU loads and the I/O parameters.

2. Hooking up Two Workstations Peer-to-peer

The second step in the installation process is to connect the two workstations with a single pair of fiber optic lines without using a switch. Figure 49 shows these two workstations in this line driver test configuration.

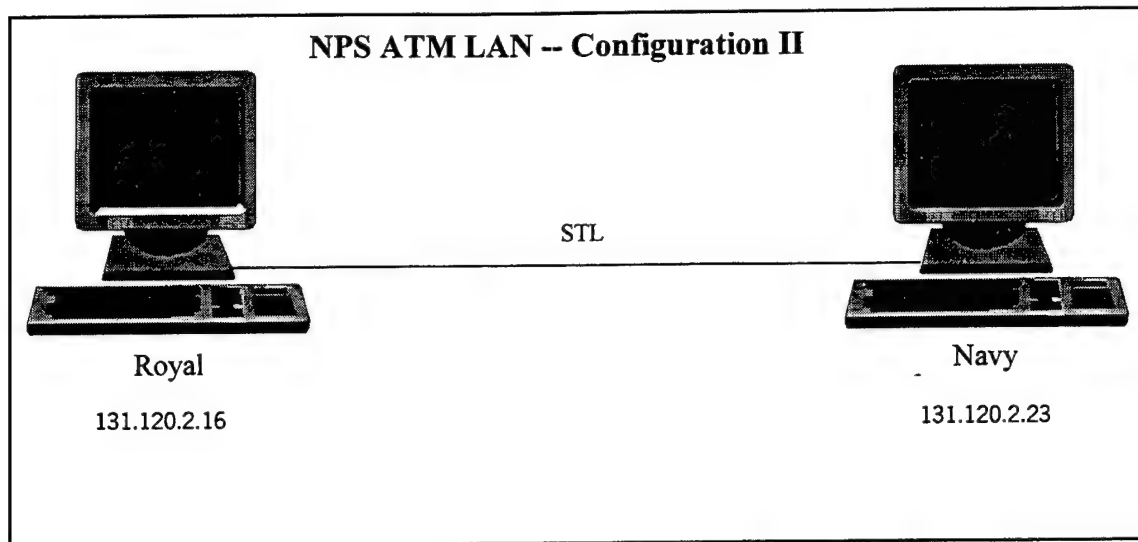


Figure 49. Peer-to-Peer.

nttcp benchmark testing is performed from Royal to Navy. The following command is issued to a Royal command window:

```
nttcp -r
```

The following command is issued to a Navy command window:

```
nttcp -t atm-navy
```

with a result of 62 Mbps with no CPU load at the time on either machine. The machine name *atm-Navy* is used so that the *nttcp* will run across the ATM link vice the Ethernet link. The Unix *ping*, *ftp*, *telnet*, and *rlogin* commands are also used satisfactorily to test out the link.

3. Hooking up a Single ATM Switch to Two Workstations

The third step in the installation process is to connect the two workstations with a single Cisco ATM A-100 HyperSwitch. Figure 50 shows the two workstations in this configuration.

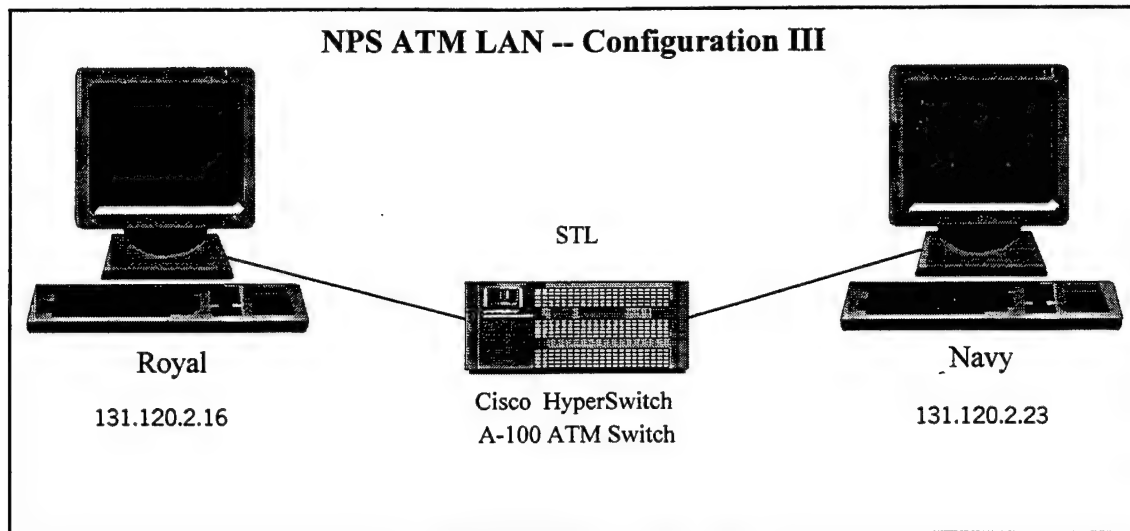


Figure 50. Single-Switched.

nttcp benchmark testing is performed from Royal to Navy. The following command is issued to a Royal command window:

```
nttcp -r
```

The following command is issued to a Navy command window:

```
nttcp -t navy
```

with a result of 52 Mbps. Due to oceanographic research in progress, there is a substantial CPU load when this measurement is done. The Unix *ping*, *ftp*, *telnet*, and *rlogin* commands are also used satisfactorily to test out the link.

4. Hooking up Two ATM Switches

The forth step in installation process is to connect the workstations with two switches, as shown in Figure 51. Cisco ATM A-100 HyperSwitches are used for both switches.

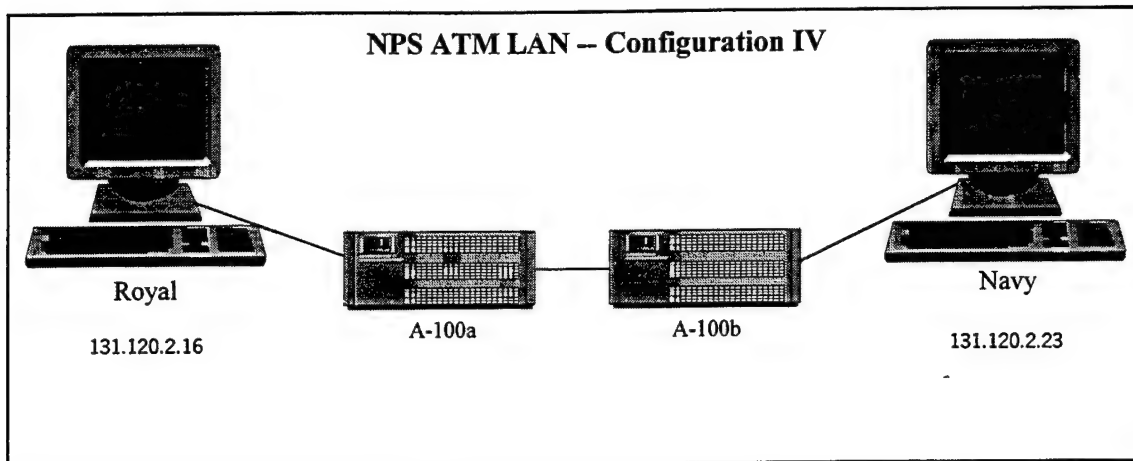


Figure 51. Dual-Switched.

nttcp benchmark testing is performed from Royal to Navy. The following command is issued to a Royal command window:

```
nttcp -r
```

The following command is issued to a Navy command window:

```
nttcp -t navy
```

with a result of 105 Mbps. There was little CPU load when this measurement is done.

The Unix *ping*, *ftp*, *telnet*, and *rlogin* commands are also used satisfactorily to test out the link.

5. Connecting to BayNet

The final step in the installation process is to connect the school's ATM LAN to CalREN and BayNet. Figure 52 displays the initial configuration. An account was established for UCSC's "Cyclone" computer, accessed across the BayNet cloud.

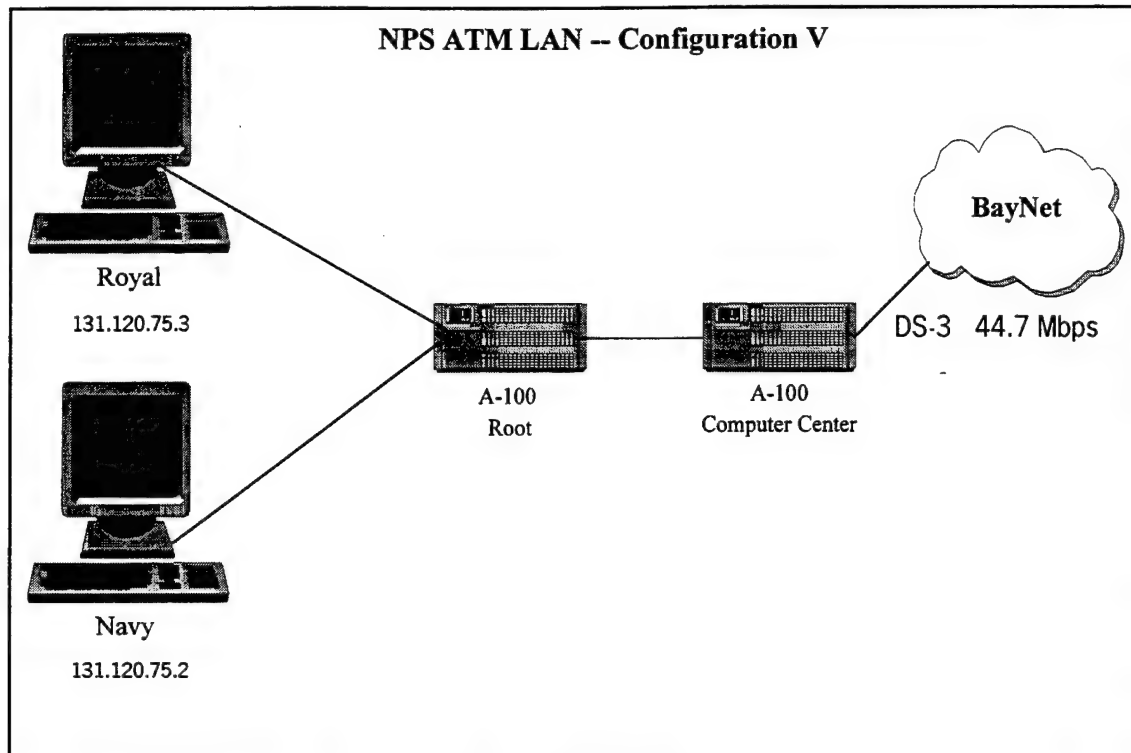


Figure 52. BayNet Configuration.

nttcp benchmark testing is performed from Royal to Cyclone. The following command is issued to a Cyclone command window:

```
nttcp -r
```

The following command is issued to a Royal command window:

```
nttcp -t cyclone
```

with a result of 8.2 Mbps. As discussed in Chapter IV, BayNet limits traffic to 10 Mbps for the PVC on this link. The Unix *ping*, *ftp*, *telnet*, and *rlogin* commands are also used satisfactorily to test out the link.

An X-terminal session is established across the BayNet link. The following command is issued to a Royal command window:

xterm +

The following command is issued to a Cyclone command window:

setenv DISPLAY royal:0.0

then the Unix X-Window commands listed in Figure 53 are successfully run across the BayNet link.

The xpsview program is a particularly useful test. This program is a postscript viewer. A postscript file located on Cyclone is opened and read across the ATM link at Royal.

- *xcalc*
- *xcalendar*
- *xclock*
- *xclipboard*
- *xeyes*
- *xfig*
- *xlogo*
- *xman*
- *xpsview*

Figure 53. X-Window Applications.

D. FAILURE TO SUSTAIN BAYNET CONNECTIVITY

After connecting to BayNet and running X-Window applications across the ATM link, IIRG's goal was to multicast a local workshop ("Web Content and Access Workshop, Monterey Bay '96") to UCSC using the ATM link. We were unable to accomplish this. The link unexpectedly went down, and the only person with ATM expertise was not available to assist in getting the link back up. This caused us not to be able to transmit the workshop to regional partners interested in participating.

We also failed in hooking up Professor Wendell Nuss or in establishing more than intermittent connectivity across the ATM link to UCSC. These failures are both technological and managerial in nature. A full discussion of these disappointments is provided in Chapter VII.

E. CONCLUSION

This chapter provides the results of experimental testing performed on the NPS ATM LAN. The testing performed is done by using standard Unix commands. The tests are performed for each of the five configurations during the LAN's evolution: stand-alone configuration, then peer-to-peer, through one switch, through two switches, and finally across campus and out to BayNet. A program provided by Fore called *nttcp* is used for most benchmark testing. The results of the testing are satisfactory during the LAN's evolution and demonstrate connectivity across all links.

VII. CONCLUSIONS AND RECOMMENDATIONS

A. INTRODUCTION

This chapter draws conclusions about ATM. Barriers to future work and recommendations are made concerning the Navy's and DoD's deployment of ATM technology in the future.

B. CONCLUSIONS

NPS has successfully implemented an ATM LAN across campus. This LAN is being used to perform research into the capabilities of ATM and to demonstrate various research initiatives at NPS using ATM technology. The strengths and weaknesses of ATM can now be tested first hand.

NPS was unsuccessful at establishing sustainable ATM connections regionally or globally. The technical and administrative reasons for that significant failure are discussed below. We believe that there are currently five fatal flaws associated with ATM that discourage widespread deployment. These flaws are listed in Figure 57. Correcting these flaws will likely make ATM connectivity as appealing as promised. Until that day, ATM network planning and deployment remain problematic. Finally, recommendations are made for future research into ATM at NPS.

C. FUTURE PLANS

1. Proposed NPS ATM LAN Extensions

The ATM LAN is currently composed of two workstations, Royal and Navy, connected by two switches to the BayNet cloud, as shown in Figure 54.

The future desired network configuration is shown in Figure 55. This reflects the short-term backbone network for the campus that connects Ingersoll (VisLab and CC), Root (STL and Professor Nuss), and Spanagel (CS) halls with BayNet. All of the internal lines will be running OC-3 (155 Mbps) rates, which is the limit of the Cisco A-100 switches. The BayNet connection is presently set to DS-3 (44.7 Mbps) by the PacBell

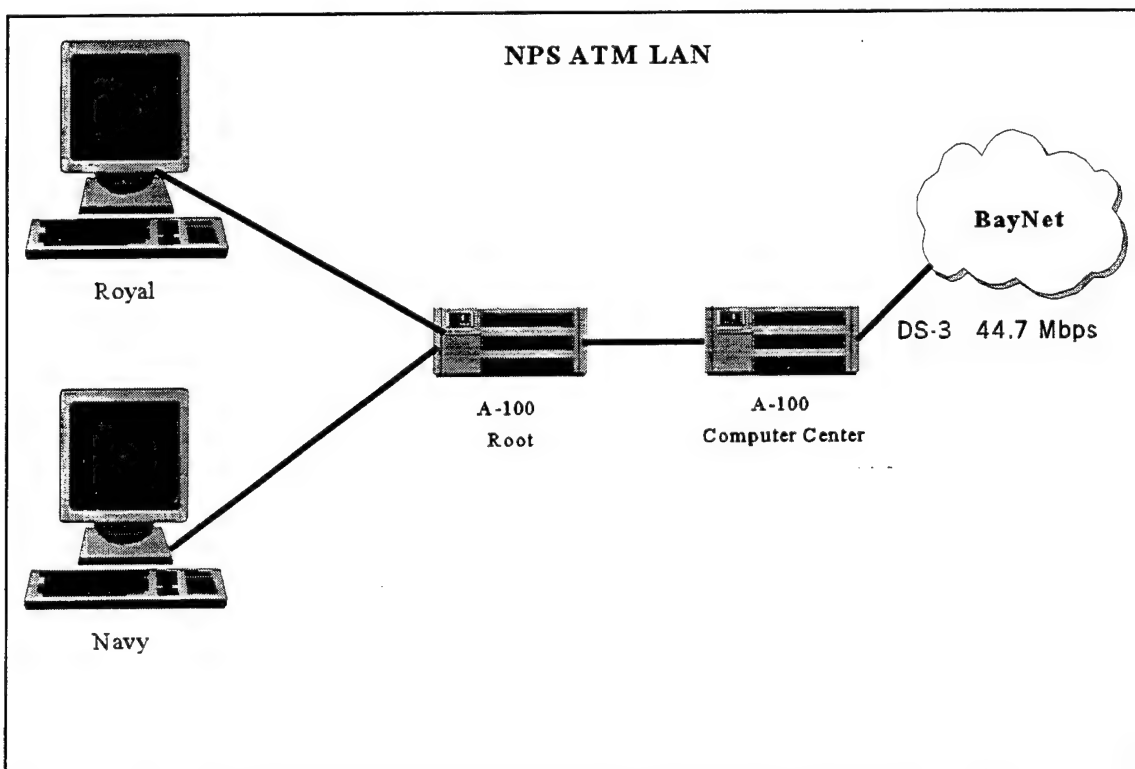


Figure 54. NPS ATM LAN Connection to BayNet Cloud.

CalREN grant, but is capable of OC-3c rates.

2. Changes in the CalREN grant

Pacific Bell's two year funding for the "free" use of their ATM service to the test-bed consortium under the CalREN program expires on 1 October 1996. Current pricing estimates by PacBell indicate that Monterey BayNet partners each may need to pay over \$30,000 annually. Our experience shows that ATM problems do not justify such expenditures, despite our long-standing enthusiasm to use ATM effectively. The five

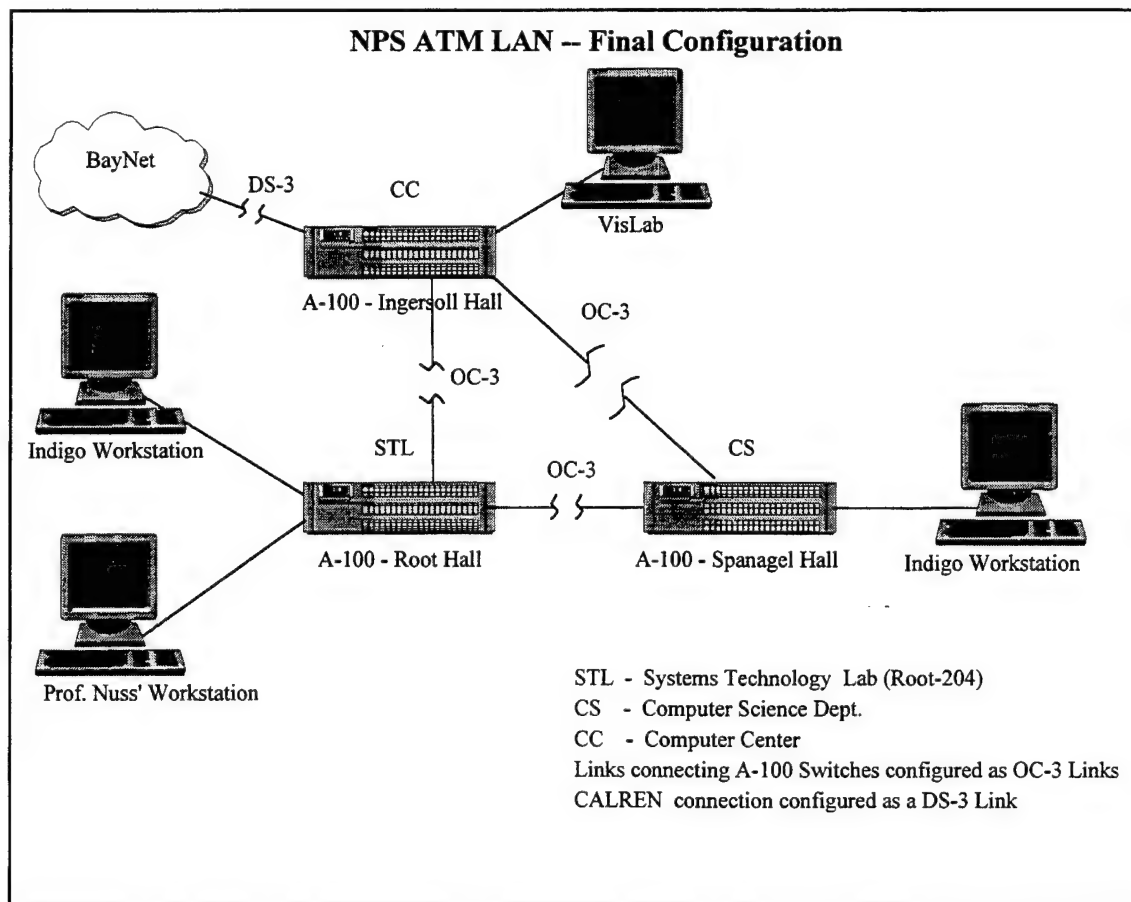


Figure 55. Final Desired Configuration, after [Advisory Group, 1995].

fatal flaws need to be addressed before significant long-term expenditures can be justified.

D. BARRIERS, FUTURE WORK, AND RECOMMENDATIONS

1. Barriers

One barrier to NPS's continuing on with ATM research external to the school is the cost of using the ATM WAN communication media provided by PacBell. PacBell fully funded the CalREN program until 1 October 1996. After this date, a 45 Mbps DS-3 access is \$2,700 per month, and a 155 Mbps OC-3c access is \$3,200 per month. The school will need to decide if the cost for ATM network access external to the school is worth the benefits of the research gained therein. It is the author's opinion that most research that NPS may want to conduct with ATM can be done "in-house" (i.e., on campus), so the need of this costly external access is unnecessary.

A second barrier is that the Cisco A-100 switches that NPS purchased are limited to 155 Mbps. This restricts the school's ability to explore the Gbps barrier, requiring investing in new switches across the campus when this technology becomes feasible. This is a legacy constraint with which we must live for the near-term future. It is quite possible that gigabit routing will be available before ATM is ready for regular internetworking.

Lastly, NPS purchased Cisco switches and Fore NICs. Therefore, the school cannot take full advantage of the proprietary benefits of either company and must set all configurations manually. UCSC did not have this constraint because their usage plans

only required proprietary Fore hardware. They and other BayNet participants were able to achieve satisfactory functionality with BayNet over a one-year period using consultant support. This is also a legacy constraint with which we must live. The lesson learned from this is that when using a technology that is still in its infancy, a single vendor solution is the best path to take. The large lesson is that widespread deployment should be avoided until switch and card interoperability problems have been demonstrated to be corrected. Proprietary restrictions need to be removed from consideration.

2. Future Work

The areas of future research are legion. Audio and video conferencing applications are bandwidth intensive and require low latency from the underlying network multicast service. IP can be run over increasingly higher capacity links to solve the bandwidth problem, but there still remains a serious latency problem with IP networks. To help compensate for this problem, much research is focusing on ATM networks. Unfortunately, it is not yet clear whether current research (at NPS and elsewhere) will be able to adequately address the five fatal flaws of ATM.

a. Multicasting

One aspect of this research is support for multicast over ATM networks. Overcoming the multicasting problem is probably the greatest problem to solve. As the Internet grows, multicasting will become more prominent, and it is a problem that needs a solution at every layer in the OSI reference model. ATM is very weak in this area since multicasting is foreign to telephone companies that only understand point-to-point,

trunked, legacy systems. On-going research has yet to find a native ATM solution to this problem. NPS wants to utilize ATM for high-bandwidth low-latency pipes which carry encapsulated IP multicast for subsequent many-to-many redistribution at endpoint IP-based LANs. Even this modest capability could not be demonstrated, which was a surprising and frustrating disappointment.

b. Reliability

One critical military feature of any DoD network is that it be highly reliable. When planning a C³I system, the redundant elements need to meet three criteria for survivability and fault tolerance. These three criteria are presented in Figure 56.

The third criteria is reliable *crossover* or *changeover*. Router-based internets automatically perform reliable crossover as routers react to congestion and broken links. Similarly, FDDI networks perform reliable crossover following single interface failure using ring wrap. As discussed in Chapter V, since ATM switches are circuit-switched, they require some signaling means to establish and reestablish a broken connection — similar to broken PSTN lines in the middle of a phone call. The

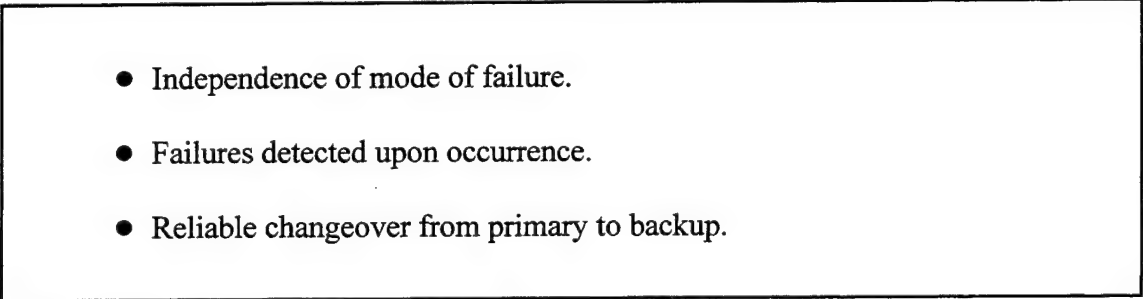
- 
- Independence of mode of failure.
 - Failures detected upon occurrence.
 - Reliable changeover from primary to backup.

Figure 56. Criteria for Survivability and Fault Tolerance, from [Buddenberg, 1995].

changeover system from primary to backup mechanisms must be reliable. It is unacceptable to have the backup systems unavailable when needed. [Buddenberg, 1995]

ATM has not solved the problem of changeover. If a connection is broken, there is no standby connection waiting to immediately take over; and this scenario is exacerbated in the already problematic multicast situation. Before DoD becomes too committed to ATM, the issues of changeover and multicast need to be explicitly and fully resolved.

c. ATM over Satellite

There has been much discussion and excitement in recent months about ATM over satellites and wireless. This area needs to be examined thoroughly as well. One of the greatest problems with ATM over satellite is not only the overhead involved with using ATM cells, but also the problems with multicasting to ships at sea and having shared data and resources. How this is accomplished needs to be thoroughly researched.

3. Recommendations for Future Work

There are some recommendations that I would make to anyone continuing on in the ATM research direction. First, that they follow the philosophy of open architecture and standardization in the system design.

Secondly, that they understand the end users' requirements and applications. Involve the users themselves in the design stage and accept their critiques during the evaluation phase; this benefited me greatly in building the LAN.

Thirdly, that they design an infrastructure which can grow with the evolution of WAN and Internet capabilities; we must get away from stove-piped solutions to our networking problems and approach the solution from a network-centric perspective.

E. RECOMMENDATIONS TO DISA AND THE U.S. NAVY

The Navy and DISA need to approach ATM cautiously and not put all their eggs in one proverbial basket. DISA has focused on ATM as being the network topology of the future. The following quote from DISA was cited in Chapter II, and it is repeated here since it shows DISA's commitment to ATM for the future of the Defense Information System Network (DISN):

[DISA] recognizes ATM technology as *the key building block* of the evolving DISN. *It alone* can satisfy the warfighter's need for the extension of high bandwidth, realtime and multi-media communications to remote theaters of operation anywhere in the world. *It alone* holds the promise of a single technology supporting high performance data, voice and video services through the seamless integration of local and wide area networks, including those in the tactical theater of operations.... Today *ATM stands alone* with its promise of solving the vexing problems of seamless sensor-to-shooter information support to the warfighter. [DISA, 1994] (emphasis mine)

Our examination and conclusion of the benefits and problems with ATM technology sounds a clarion call, warning DISA to evaluate the risk of transferring to a cutting-edge technology such as ATM. In my opinion ATM is not yet ready for prime-time, especially for national defense. Specifically, the following four "show stoppers" exist, as detailed in Figure 57.

- Interoperability between switches
- Incompatibility with IP
- Inflexibility to change
- Individuals
- Crossover

Figure 57. ATM "Show-stoppers."

1. Interoperability between switches

There is no way to guarantee communication between switches. This is seen in the communication problem encountered between the PacBell and Sprint switches. The PacBell switches were able to communicate with each other using SVCs, but the Sprint switch could only handle PVCs.

2. Incompatibility with IP

There is no native way to multicast with ATM. There is a lot of effort going into solving the IP and multicasting problems, but so far no one has come up with an acceptable solution.

3. Inflexibility to change

Myriad long-haul problems exist. These problems are especially difficult due to the change in the regulations concerning RBOC's being long-haul carriers. This includes setup and tariff issues.

4. Individuals

The human engineering problem was almost insurmountable. The "expertise" that exists in the ATM field is nominal due to the immaturity of the technology. No one at NPS had ever dealt with ATM prior to our purchasing the switches. All the expertise that we had was learned through trying to establish the LAN. Trying to get assistance is next to impossible due to the fact that so few people have any proficiency in ATM.

Some regional partners had greater but limited success. The Monterey Bay Aquarium and the San Jose Technological Museum for Innovation were able to get ATM running between their two institutions by throwing a lot of consulting money at the problem and by running strictly a single application (audio/video) only across a PVC connection. UCSC and UCSC-Ext were able to communicate with SVCs because they did not have to use Sprint's ATM switch.

5. Crossover

As discussed above, if a connection is broken, there is no standby connection waiting to immediately take over. This scenario is heightened in the already problematic multicast situation.

F. SUMMARY

The advantages of ATM are many: delay, jitter, capacity, reliability, and priority are well controlled. The network is easy to construct, but difficult to configure and run. ATM is also still costly, and seamless interoperability between vendor networks simply does not exist. If one uses the same vendor for a solution, there might be better, quicker,

and easier results than by mixing vendor's equipment — like NPS did. DoD, however, cannot adopt a proprietary networking technology that may ultimately constrain interoperability, sacrifice critical military features, or increase cost without bound. DoD must have open standards-based solutions as the long term goal.

The ability to build and run an ATM LAN have been proven at NPS. The other problems with ATM still remain, however. As with any new technology, there are risks. All of the issues discussed above strongly indicate that the Navy and DoD need take a cautious approach to the introduction of ATM technology. ATM is just one technology in the mix. Like every link-layer technology, there will be lots of ideas about how to use it better. The market and the IP standards world will have to sort out which ones are good or bad. In the long-term, ATM's inability to handle arbitrarily large, many-to-many multicast may prove to be its Achilles heel. More work is necessary before committing to wide-scale deployment of ATM (e.g., Navy- or DoD-wide).

The CalREN ATM grant was a failure for NPS. There were no shared goals or shared time-line for CalREN, and therefore the coordination and objectives between CalREN members were not congruent. These led to ATM being a social and administrative failure regionally. This is regrettable.

ATM was also a technical failure. We were able to get ATM working across campus, but trying to communicate with regional or global participants was impossible. The limited success that NPS had in communicating with UCSC was limited to running XTerminal sessions across the ATM link — not a very robust technical achievement for a

10 Mbps pipe between the two institutions that will cost \$3,000 per month starting October 1, 1996. We do not expect that other institutions will be willing to follow our example and work with ATM on an intensive daily basis, only to achieve unacceptable, flawed results.

ATM demonstrated considerable weakness with interoperability. Trying to get long-haul providers technically passing cells to each other was impossible since Sprint was unable to use SVCs. UCSC had limited success in communicating with their UCSC Extension campus since they only had to communicate with one carrier (PacBell). NPS, having to communicate with two carriers to reach our nearest neighbor, was unable to make ATM work successfully. We observed repeated examples of this flaw on a national basis, most notably with the I-WAY project. The I-WAY demonstrated impressive and achievable technical successes which nevertheless failed to convince telecommunications providers to offer ATM connectivity between supercomputer centers at less than astronomical prices.

There are also problems with the changing legislation of the Telecommunications Act of 1996. This allows the RBOC to provide long-haul carrier service. This will cause carrier services to be in flux as to what they will provide, how it will be provided, and what interoperability problems will exist. Despite the Telecommunication Act and FCC rulings, telcos are not ready to overcome long-haul problems (such as LATA boundaries). This demonstrated inertia in the face of statutory change is troubling.

ATM is worthwhile if the five major flaws — interoperability between switches, incompatibility with IP, inflexibility to change, individuals, and crossover — are corrected. These issues need to be rectified before the Navy or DoD commits to wide-scale deployment.

LIST OF REFERENCES

Alles, Anthony, *ATM Internetworking*, Cisco Systems, Inc., San Jose, Calif., May 1995.
Available at <http://www.cisco.com/warp/public/614/12.html>

ATDNet, *ATDnet Experiment Plan*, Dec., 1994. Available at
<http://www.atd.net/ATDNET/experiment.html>

Ballistic Research Lab, *nttcp.*, 1984. Available at <http://ftp.arl.mil/ftp/pub/ttcp/README>

Bigelow, R. Jon, *Internetworking: Planning and Implementing a Wide Area Network (WAN) for K-12 Schools*, Master's Thesis, Naval Postgraduate School, Monterey, California, 1995.

Black, U., *Data Networks, Concepts, Theory and Practice*, Englewood Cliffs, New Jersey, Prentice Hall, 1989.

_____, *Physical Level Interfaces and Protocols*, Los Alamitos, California, IEEE Computer Society Press, 1988.

Buddenburg, Rex, *High Availability*, Monterey, California, 1995. Available at
http://dubhe.cc.nps.navy.mil/~budden/book/two.avail_.html

CalREN, Project: #ATMN-014 *ATM As An Enabling Technology for Tele-education, Telescience, and Electronic Libraries linking the Silicon Valley, Santa Cruz, and the Monterey Peninsula*, 1994. Available at
<http://www.pacbell.com/SuperHi/CalREN/Projects/atmnorth-7.html>

Case, J.D., et al, *Network Management and the Design of SNMP*, ConneXions, The Interoperability Report, Vol. 3, March 1989.

_____, *Introduction to the Simple Gateway Monitoring Protocol*, IEEE Network, March 1988.

Cisco LightStream 100 User Guide, Cisco Systems, Inc., San Jose, Calif., 1995.
Available at
http://www.cisco.com/univercd/data/doc/hardware/fusion/a100/ls100_ug.htm

Clark, W., *SNA Internetworking*. ConneXions, The Interoperability Report, Vol. 6, No. 3, March 1992.

Coltun, R., *OSPF, An Internet Routing Protocol*, ConneXions, The Interoperability Report, Vol. 3, No. 8, August 1989.

Comer, D.E., *Internetworking with TCP/IP, Principles, Protocols, and Architecture, Vol. I*, 2nd Ed. Englewood Cliffs, New Jersey, Prentice Hall, 1991.

Crawley, Eric S., *Multicast Routing over ATM*, Internet-draft, Feb. 22, 1996, expires Aug. 22, 1996. Available at <http://www.es.net/pub/internet-drafts/draft-crawley-mcast-rout-over-atm-00.txt>

Davidson, J., *An Introduction to TCP/IP*, New York, New York, Springer-Verlag, 1992.

Dennis, Ronald, *Internetworking: IP/ATM LAN Security*, Master's Thesis, Naval Postgraduate School, Monterey, California, September 1996. Available at <http://www.stl.nps.navy.mil/~iirg/>

DISA, *Asynchronous Transfer Mode (ATM) Technology Initiatives in the Department of Defense (DoD)*, 1994. Available at <http://disa11.disa.atd.net/atmdemo/atmdemo.html>

Ebrahim, Zahir, *A Brief Tutorial on ATM*, March 5, 1992. Available at <http://juggler.lanl.gov/lanp/atm.tutorial.html>

Edwards, Evan, *Internetworking: Network Monitoring Aspects of High Bandwidth, Low Latency Global Networks*, Master's Thesis, Naval Postgraduate School, Monterey, California, September 1996. Available at <http://www.stl.nps.navy.mil/~iirg/>

Emswiler, Tracy L., *Distance Learning Over the Multicast Backbone (MBone)*, Master's Thesis, Naval Postgraduate School, Monterey, California, September 1995.

Erdogan, Ridvan, *Internetworking: Using Global ATM Networks for Live Multicast Audio/Video Distribution*, Master's Thesis, Naval Postgraduate School, Monterey, California, September 1995. Available at <http://www.stl.nps.navy.mil/~iirg/>

Feibel, Werner, *The Encyclopedia of Networking*, Second Edition, Sybex Network Press, Alameda, California, 1995.

ForeRunner VMA-200 ATM VMEbus Adapter for Silicon Graphics, Fore Systems, Inc., Warrendale, PA, 1995. Available at <http://www.fore.com/html/products/adpt/vma200a.html>

Forum, *The ATM Forum Consortium Home Page*. Available at <http://www.atmforum.com/>

Freeman, Roger L., *Telecommunication Transmission Handbook*, John Wiley & Sons, New York 1991.

Garcia-Luna-Aceves, J.J., *Loop-Free Routing Using Diffusing Computations*, IEEE/ACM, Transactions on Networking, Vol. 1, No. 1, 1993.

Green, J.K., *Telecommunications*, 2nd Ed. Homewood, Illinois, Business One Irwin, 1992.

Hagans, R., *Components of OSI, ES-IS Routing*, ConneXions, The Interoperability Report, Vol. 3, No. 8, August 1989.

Hares, S., *Components of OSI, Inter-Domain Routing Protocol (IDRP)*, ConneXions, The Interoperability Report, Vol. 6, No. 5, May 1992.

Interview between the NPS Networking Advisory Group and author, Nov. 22, 1995

Interview between Arul Ananthanarayanan, Senior UNIX Systems and ATM Manager, UCSC, and author, Mar. 7, 1995.

Interview between Dave Norman, Manager of NPS Computer Systems Support, and author, Jan. 16, 1996.

Interview between Roy Romo, Assistant Manager of NPS Computer Systems Support, and author, Dec. 15, 1996.

I-WAY, *The Information Wide Area Year*, 1995. Available at <http://www.iway.org/>

Jarrin, K.M., *Beyond IPv6 and Coexistence with ATM*. Available at <http://www.ece.arizona.edu:8001/~kmjarrin/homework/hw2.html#beyond>

Joyce, S.T., and Walker II, J.Q., *Advanced Peer-to-Peer Networking (APPN), An Overview*, ConneXions, The Interoperability Report, Vol. 6, No. 10, October 1992.

- Kousky, K., *Bridging the Network Gap*, LAN Technology, Vol. 6, No. 1, January 1990.
- Leahy, Kevin M., *The Proposed Naval Postgraduate School Campus Network: Computer Communications for the 1990's*, Master's Thesis, Naval Postgraduate School, Monterey, California, March 1988.
- Leinwand, A., and Fang, K., *Network Management, A Practical Perspective*, Reading, Massachusetts, Addison-Wesley Publishing Company, Inc., 1993.
- Lippis, N., *The Internetwork Decade*, Data Communications, Vol. 20, No. 14, October 1991.
- Martin, J., *SNA, IBM's Networking Solution*, Englewood Cliffs, New Jersey, Prentice Hall, 1987.
- _____, *Local Area Networks, Architectures and Implementations*, Englewood Cliffs, New Jersey, Prentice Hall, 1989.
- McNamara, J.E., *Local Area Networks*. Digital Press, Educational Services, Digital Equipment Corporation, 12 Crosby Drive, Bedford, MA 01730, n.d.
- Medin, M., *The Great IGP Debate---Part Two, The Open Shortest Path First (OSPF) Routing Protocol*, ConneXions, The Interoperability Report, Vol. 5, No. 10, October 1991.
- Meijer, A., *Systems Network Architecture, A Tutorial*, New York, New York, John Wiley & Sons, Inc., 1987.
- Miller, M.A., *LAN Protocol Handbook*, San Mateo, California, M&T Books, 1990.
- NRL Network Research Navigator*, August 1996. Available at <http://netlab.itd.nrl.navy.mil/atm.html>
- NSSN Architecture Connectivity Solutions*, ComPath Technologies, Inc., Niantic CT., 1995. Available at <http://www.com-path.com/ATM/NSSN/NSSN.htm>
- O'Reilly, T., and Todino G., *Managing UUCP and Usenet*, 10th Ed., Sebastopol, California, O'Reilly & Associates, Inc., 1992.
- Partridge, Craig, *Gigabit Networking*, Second Edition, Addison-Wesley, New York, 1994.

Perlman, R., *Interconnections, Bridges and Routers*, Reading, Massachusetts, Addison-Wesley Publishing Company, Inc., 1992.

Perlman, R., and Callon R., *The Great IGP Debate---Part One, IS-IS and Integrated Routing*, ConneXions, The Interoperability Report, Vol. 5, No. 10, October 1991.

Rettinger, Leigh A., *Desktop Videoconferencing: Technology and Use for Remote Seminar Delivery*, Master's Thesis, North Carolina State University, Raleigh, North Carolina, July 1995. Available at http://www2.ncsu.edu/eos/service/ece/project/succeed_info/larettin/thesis/abs.html

Rose, M.T., *The Open Book, A Practical Perspective on OSI*, Englewood Cliffs, New Jersey, Prentice Hall, 1990.

_____, *The Simple Book, An Introduction to Management of TCP/IP-based Internets*, Englewood Cliffs, New Jersey, Prentice Hall, 1991.

Ross, F.E., *FDDI---A Tutorial*, IEEE Communications Magazine, Vol. 24, No. 5, May 1986.

Sapronov, Walter, and Franklin, Anne E., 1996, *Summary of the Telecommunications Act of 1996*. Available at <http://www.brp.com/netline/corner/ws0396.html>

Schlar, S.K., *Inside X.25, A Manager's Guide*, New York, New York, McGraw-Hill, Inc., 1990.

Schwartz, M., *Telecommunications Networks, Protocols, Modeling, and Analysis*, Reading, Massachusetts, Addison-Wesley Publishing Company, Inc., 1987.

Sherman, K., *Data Communications, A User's Guide*, Englewood Cliffs, New Jersey, Prentice Hall, 1990.

Siu, Kai-Yeung, and Jain, Raj, *A Brief Overview of ATM: Protocol Layers, LAN Emulation, and Traffic Management*, 1995. Available at <http://www.cis.ohio-state.edu/~jain/paper.html>

Spragins, J.D., et al., *Telecommunications Protocols and Design*, Reading, Massachusetts, Addison-Wesley Publishing Company, Inc., 1991.

Stallings, W., *Data and Computer Communications*, New York, New York, Macmillan Publishing Company, 1991.

_____, *Handbook of Computer-Communications Standards, Vols. 1-3*, Carmel, Indiana, Howard W. Sams, Inc., 1990.

_____, *Local Networks*, 3rd Ed. New York, New York, Macmillan Publishing Company, 1990.

Sunshine, C.A. (Ed.), *Computer Network Architectures and Protocols*, 2nd Ed. New York, New York, Plenum Press, 1989.

Suzuki, T., *ATM Adaptation Layer Protocol*, IEEE Communication Magazine, Vol. 32, No. 4, Apr. 1994.

Tamer, Murat, *Internetworking: Multicast and ATM Network Prerequisites for Distance Learning*, Master's Thesis, Naval Postgraduate School, Monterey, California, September 1996. Available at <http://www.stl.nps.navy.mil/~iirg/>

Tannenbaum, A.S., *Computer Networks*, 3rd Edition, Englewood Cliffs, New Jersey, Prentice Hall, 1995.

Terplan, K., *Communication Networks Management*. Englewood Cliffs, New Jersey, Prentice Hall, 1992.

Tiddy, Michael, *Internetworking: Economical Storage and Retrieval of Digital Audio and Video for Distance Learning*, Master's Thesis, Naval Postgraduate School, Monterey, California, September 1996. Available at <http://www.stl.nps.navy.mil/~iirg/>

Tsuchiya, P., *Components of OSI, IS-IS Intra-Domain Routing*, ConneXions, The Interoperability Report, Vol. 3, No. 8, August 1989.

_____, *Components of OSI, Routing (An Overview)*, ConneXions, The Interoperability Report, Vol. 3, No. 8, August 1989.

Varma, Anujan, *Introduction to ATM Networks*, Lecture Notes and Course Video, University of California Santa Cruz, 1995.

Wheless, Glen H., *The Chesapeake Bay Virtual Ecosystem (CBVE): Interacting with a Coupled Bio-physical Simulation*. Available at <http://www.ccpo.odu.edu/~wheless/SC95.html>

_____, Eileen E. Hofmann, John M. Klinck, Cathy M. Lascara, Arnoldo Valle-Levinson, Donald P. Brutzman, and Bill Sherman, *Virtual Chesapeake Bay*:

Interacting with a Coupled Physical/ Biological Model, IEEE Computer Graphics and Applications, Vol. 16, No. 4, July 1996.

Wiedenhoeft, Paul E., *Analysis of the Naval Postgraduate School Computer Network Architecture*, Master's Thesis, Naval Postgraduate School, Monterey, California, September 1994.

Yao, Engui, *ATM -- The New Technology for Tomorrow's B-ISDN*. Available at <http://www2.msstate.edu/~eyl/paper/tkt8793paper.html>

Zimmerman, H., *OSI Reference Model---The ISO Model of Architecture for Open Systems Interconnection*, IEEE Transactions on Communications COM-28, No. 4, April 1980.

APPENDIX A. ABBREVIATIONS AND DEFINITIONS

The following listing of abbreviations and definitions is an abridged and extended version of Cisco's *Internetworking Terms and Acronyms*. The abridgement is based on terms used in this thesis. The entire glossary can be found at <http://www.erl.noaa.gov/noc/cisco/data/doc/cintrnet/ita.htm>

4B/5B local fiber: 4-byte/5-byte local fiber. Fiber channel physical media used for FDDI and ATM. Supports speeds of up to 100 Mbps over multimode fiber.

8B/10B local fiber: 8-byte/10-byte local fiber. Fiber channel physical media that supports speeds up to 149.76 Mbps over multimode fiber.

AAL: ATM adaptation layer. Service-dependent sublayer of the data link layer. The AAL accepts data from different applications and presents it to the ATM layer in the form of 48-byte ATM payload segments. AALs consist of two sublayers, CS and SAR. AALs differ on the basis of the source-destination timing used, whether they use CBR or VBR, and whether they are used for connection-oriented or connectionless mode data transfer. At present, the four types of AAL recommended by the ITU-T are AAL1, AAL2, AAL3/4, and AAL5.

AAL1: ATM adaptation layer 1. One of four AALs recommended by the ITU-T. AAL1 is used for connection-oriented, delay-sensitive services requiring constant bit rates, such as uncompressed video and other isochronous traffic.

AAL2: ATM adaptation layer 2. One of four AALs recommended by the ITU-T. AAL2 is used for connection-oriented services that support a variable bit rate, such as some isochronous video and voice traffic.

AAL3/4: ATM adaptation layer 3/4. One of four AALs (merged from two initially distinct adaptation layers) recommended by the ITU-T. AAL3/4 supports both connectionless and connection-oriented links, but is primarily used for the transmission of SMDS packets over ATM networks.

AAL5: ATM adaptation layer 5. One of four AALs recommended by the ITU-T. AAL5 supports connection-oriented, VBR services, and is used predominantly for the transfer of classical IP over ATM and LANE traffic. AAL5 uses SEAL and is the least complex of the current AAL recommendations. It offers low bandwidth overhead and simpler processing requirements in exchange for reduced bandwidth capacity and error-recovery capability.

ABR: available bit rate. QoS class defined by the ATM Forum for ATM networks. ABR is used for connections that do not require timing relationships between source and destination. ABR provides no guarantees in terms of cell loss or delay, providing only best-effort service. Traffic sources adjust their transmission rate in response to information they receive describing the status of the network and its capability to successfully deliver data.

ACK: acknowledgment. Notification sent from one network device to another to acknowledge that some event (for example, receipt of a message) has occurred.

ACR: allowed cell rate. Parameter defined by the ATM Forum for ATM traffic management. ACR varies between the MCR and the PCR, and is dynamically controlled using congestion control mechanisms.

ADPCM: adaptive differential pulse code modulation. Process by which analog voice samples are encoded into high-quality digital signals.

ADSU: ATM DSU. Terminal adapter used to access an ATM network via an HSSI-compatible device.

ANSI: American National Standards Institute. Voluntary organization comprised of corporate, government, and other members that coordinates standards-related activities, approves U.S. national standards, and develops positions for the United States in international standards organizations. ANSI helps develop international and U.S. standards relating to (among other things) communications and networking. ANSI is a member of the IEC and the ISO.

ARP: Address Resolution Protocol. Internet protocol used to map an IP address to a MAC address.

ASCII: American Standard Code for Information Interchange. 8-bit code for character representation (7 bits plus parity).

ATDM: asynchronous time-division multiplexing. Method of sending information that resembles normal TDM, except that time slots are allocated as needed rather than preassigned to specific transmitters.

ATM: Asynchronous Transfer Mode. International standard for cell relay in which multiple service types (such as voice, video, or data) are conveyed in fixed-length (53-byte) cells. Fixed-length cells allow cell processing to occur in hardware, thereby reducing transit delays. ATM is designed to take advantage of high-speed transmission media such as E3, SONET, and T3.

ATM Forum: International organization jointly founded in 1991 by Cisco Systems, NET/ADAPTIVE, Northern Telecom, and Sprint that develops and promotes standards-based implementation agreements for ATM technology. The ATM Forum expands on official standards developed by ANSI and ITU-T, and develops implementation agreements in advance of official standards.

ATM layer: Service-independent sublayer of the data link layer in an ATM network. The ATM layer receives the 48-byte payload segments from the AAL and attaches a 5-byte header to each, producing standard 53-byte ATM cells. These cells are passed to the physical layer for transmission across the physical medium.

ATMM: ATM management. Process that runs on an ATM switch that controls VCI translation and rate enforcement.

ATM user-user connection: Connection created by the ATM layer to provide communication between two or more ATM service users, such as ATMM processes. Such communication can be unidirectional (using one VCC) or bidirectional (using two VCCs).

BARRNet: Bay Area Regional Research Network. Regional IP-based routed network serving the San Francisco Bay Area. The BARRNet backbone is composed of four University of California campuses (Berkeley, Davis, Santa Cruz, and San Francisco), Stanford University, Lawrence Livermore National Laboratory, and NASA Ames Research Center. BARRNET is now part of BBN Planet.

BER: bit error rate. The ratio of received bits that contain errors to total received bits.

BISDN: Broadband ISDN. ITU-T communication standards designed to handle high-bandwidth applications such as video. BISDN currently uses ATM technology over SONET-based transmission circuits to provide data rates from 155 to 622 Mbps and beyond.

BOOTP: IP/UDP bootstrap protocol (BOOTP) which allows a diskless client machine to discover its own IP address, the address of a server host, and the name of a file to be loaded into memory and executed.

BT: burst tolerance. Parameter defined by the ATM Forum for ATM traffic management. For VBR connections, BT determines the size of the maximum burst of contiguous cells that can be transmitted.

CAC: call admission control. Traffic management mechanism used in ATM networks that determines whether the network can offer a path with sufficient bandwidth for a requested VCC.

Category 5 cabling: One of five grades of UTP cabling described in the EIA/TIA-586 standard. Category 5 cabling is used for running CDDI and can transmit data at speeds up to 100 Mbps.

CBR: constant bit rate. QoS class defined by the ATM Forum for ATM networks. CBR is used for connections that depend on precise clocking to ensure undistorted delivery.

CCS: common channel signaling. Signaling system used in telephone networks that separates signaling information from user data. A specified channel is exclusively designated to carry signaling information for all other channels in the system.

CDDI: copper distributed data interface. Implementation of FDDI protocols over STP and UTP cabling. CDDI transmits over relatively short distances (about 100 meters), providing data rates of 100 Mbps using a dual-ring architecture to provide redundancy.

CDVT: cell delay variation tolerance. Parameter defined by the ATM Forum for ATM traffic management. In CBR transmissions, determines the level of jitter that is tolerable for the data samples taken by the PCR.

CIA: classical IP over ATM. Specification for running IP over ATM in a manner that takes full advantage of the features of ATM. The following RFCs and I-Ds address the specifications and controversy:

<http://www2.es.net/pub/internet-drafts/draft-armitage-ipatm-encaps-02.txt> — Issues surrounding a new encapsulation for IP over ATM

<http://www2.es.net/pub/internet-drafts/draft-ietf-ipatm-arch-00.txt> — IP Multicasting over ATM: System Architecture Issues

<http://www2.es.net/pub/internet-drafts/draft-ietf-ipatm-arequipa-00.txt> — Application REQUested IP over ATM (AREQUIPA)

<http://www2.es.net/pub/internet-drafts/draft-ietf-ipatm-bcast-03.txt> — IP Broadcast over ATM Networks

<http://www2.es.net/pub/internet-drafts/draft-ietf-ipatm-classic2-02.txt> — Classical IP and ARP over ATM

<http://www2.es.net/pub/internet-drafts/draft-ietf-ipatm-framework-doc-07.txt> — IP over ATM: A Framework Document

<http://www2.es.net/pub/internet-drafts/draft-ietf-ipatm-ipmc-12.txt> — Support for Multicast over UNI 3.0/3.1 based ATM Networks

<http://www2.es.net/pub/internet-drafts/draft-ietf-ipatm-ipv6nd-02.txt> — IPv6 and Neighbor Discovery over ATM

<http://www2.es.net/pub/internet-drafts/draft-ietf-ipatm-mib-02.txt> — Definitions of Managed Objects for Classical IP and ARP Over ATM Using SMIPv2

<http://www2.es.net/pub/internet-drafts/draft-ietf-ipatm-sig-uni-v40-00.txt> — ATM Signaling Support for IP over ATM - UNI 4.0 Update

<http://andrew2.andrew.cmu.edu/rfc/rfc1626.html> — Default IP MTU for use over ATM AAL5

<ftp://ftp.com21.com/pub/ip-atm/rfc/rfc1755.txt> — ATM Signaling Support for IP over ATM

<ftp://ftp.com21.com/pub/ip-atm/rfc/rfc1754.txt> — IP over ATM Working Group's - Recommendations for the ATM Forum's Multiprotocol BOF - Version 1

<ftp://ftp.com21.com/pub/ip-atm/rfc/rfc1626.txt> — Default IP MTU for use over ATM AAL5

<ftp://ftp.com21.com/pub/ip-atm/rfc/rfc1577.txt> — Classical IP and ARP over ATM

<ftp://ftp.com21.com/pub/ip-atm/rfc/rfc1483.txt> — Multiprotocol Encapsulation over ATM Adaptation Layer 5

CLP: cell loss priority. Field in the ATM cell header that determines the probability of a cell being dropped if the network becomes congested. Cells with CLP = 0 are insured traffic, which is unlikely to be dropped. Cells with CLP = 1 are best-effort traffic, which might be dropped in congested conditions in order to free up resources to handle insured traffic.

CoS: class of service. Indication of how an upper-layer protocol requires that a lower-layer protocol treat its messages. In SNA subarea routing, CoS definitions are used by subarea nodes to determine the optimal route to establish a given session. A CoS definition comprises a virtual route number and a transmission priority field.

CPCS: common part convergence sublayer. One of the two sublayers of any AAL. The CPCS is service-independent and is further divided into the CS and the SAR sublayers. The CPCS is responsible for preparing data for transport across the ATM network, including the creation of the 48-byte payload cells that are passed to the ATM layer.

CRC: cyclic redundancy check. Error-checking technique in which the frame recipient calculates a remainder by dividing frame contents by a prime binary divisor and compares the calculated remainder to a value stored in the frame by the sending node. Detects most cell errors.

CS: convergence sublayer. One of the two sublayers of the AAL CPCS, responsible for padding and error checking. PDUs passed from the SSCS are appended with an 8-byte trailer (for error checking and other control information) and padded, if necessary, so that the length of the resulting PDU is divisible by 48. These PDUs are then passed to the SAR sublayer of the CPCS for further processing.

DCC: Data Country Code. One of two ATM address formats developed by the ATM Forum for use by private networks. Adapted from the subnetwork model of addressing in which the ATM layer is responsible for mapping network layer addresses to ATM addresses.

D channel: data channel. Full-duplex, 16-kbps (BRI) or 64-kbps (PRI) ISDN channel.

DE: discard eligible traffic. ATM cells that have their CLP bit set to 1. If the network is congested, tagged traffic can be dropped to preferentially deliver higher-priority traffic. Sometimes called tagged traffic.

DQDB: Distributed Queue Dual Bus. Data link layer communication protocol, specified in the IEEE 802.6 standard, designed for use in MANs. DQDB, which permits multiple systems to interconnect using two unidirectional logical buses, is an open standard that is designed for compatibility with carrier transmission standards, and is aligned with emerging standards for BISDN.

- DHCP: Dynamic Host Configuration Protocol (DHCP). Provides a mechanism for transmitting configuration parameters to hosts using the TCP/IP protocol suite. The format of DHCP messages is based on the format of BOOTP messages, so that, in certain circumstances, DHCP and BOOTP participants may exchange messages.
- DS-0: digital signal level 0. Framing specification used in transmitting digital signals over a single channel at 64-kbps on a T1 facility.
- DS-1: digital signal level 1. Framing specification used in transmitting digital signals at 1.544-Mbps on a T1 facility (in the United States) or at 2.108-Mbps on an E1 facility (in Europe).
- DS-1/DTI: DS-1 domestic trunk interface. Interface circuit used for DS-1 applications with 24 trunks.
- DS-3: digital signal level 3. Framing specification used for transmitting digital signals at 44.736-Mbps on a T3 facility.
- DSAP: destination service access point. The SAP of the network node designated in the Destination field of a packet.
- DSP: domain specific part. The part of a CLNS address that contains an area identifier, a station identifier, and a selector byte.
- DXI: Data Exchange Interface. ATM Forum specification that defines how a network device such as a bridge, router, or hub can effectively act as an FEP to an ATM network by interfacing with a special DSU that performs packet segmentation and reassembly.
- E1: Wide-area digital transmission scheme used predominantly in Europe that carries data at a rate of 2.048 Mbps. E1 lines can be leased for private use from common carriers.
- E.164: ITU-T recommendation for international telecommunication numbering, especially in ISDN, BISDN, and SMDS. An evolution of standard telephone numbers.
- E3: Wide-area digital transmission scheme used predominantly in Europe that carries data at a rate of 34.368 Mbps. E3 lines can be leased for private use from common carriers.

ELAN: emulated LAN. ATM network in which an Ethernet or Token Ring LAN is emulated using a client-server model. ELANs are composed of an LEC, an LES, a BUS, and an LECS. Multiple ELANs can exist simultaneously on a single ATM network. ELANs are defined by the LANE specification.

EoT: end of transmission. Generally, a character that signifies the end of a logical group of characters or bits.

ETSI: European Telecommunication Standards Institute. Organization created by the European PTTs and the European Community (EC) to propose telecommunications standards for Europe.

FTP: File Transfer Protocol. Application protocol, part of the TCP/IP protocol stack, used for transferring files between network nodes.

G.804: ITU-T framing standard that defines the mapping of ATM cells into the physical medium.

Gbps: gigabits per second.

ICD: International Code Designator. One of two ATM address formats developed by the ATM Forum for use by private networks. Adapted from the subnetwork model of addressing in which the ATM layer is responsible for mapping network layer addresses to ATM addresses.

IDI: initial domain identifier. In OSI, the portion of the NSAP that specifies the domain.

IEEE: Institute of Electrical and Electronics Engineers. Professional organization whose activities include the development of communications and network standards. IEEE LAN standards are the predominant LAN standards today.

IETF: Internet Engineering Task Force. Task force consisting of over 80 working groups responsible for developing Internet standards. The IETF operates under the auspices of ISOC.

IITA: Information Infrastructure Technology and Applications. Component of the HPCC program intended to ensure U.S. leadership in the development of advanced information technologies.

ILMI: Interim Local Management Interface. Specification developed by the ATM Forum for incorporating network-management capabilities into the ATM UNI.

Internet: Term used to refer to the largest global internetwork, connecting tens of thousands of networks worldwide and having a "culture" that focuses on research and standardization based on real-life use. Many leading-edge network technologies come from the Internet community. The Internet evolved in part from ARPANET. At one time, called the DARPA Internet. Not to be confused with the term internet.

internet: Short for internetwork. While an internet is a network that uses the same technology (such as TCP/IP), the term "internet" is usually used to refer to a collection of such networks interconnected with routers. Not to be confused with the Internet.

Inverse ARP: Inverse Address Resolution Protocol. Method of building dynamic routes in a network. Allows an access server to discover the network address of a device associated with a virtual circuit.

IP: Internet Protocol. Network layer protocol in the TCP/IP stack offering a connectionless-routed internetwork service. IP provides features for addressing, type-of-service specification, fragmentation and reassembly, and security.

IP address: 32-bit address (IPv4 or 128-bit address in IPv6) assigned to hosts using TCP/IP. An IP address belongs to one of five classes (A, B, C, D, or E) and is written as 4 octets separated with periods (dotted decimal format). Each address consists of a network number, an optional subnetwork number, and a host number. The network and subnetwork numbers together are used for routing, while the host number is used to address an individual host within the network or subnetwork. A subnet mask is used to extract network and subnetwork information from the IP address. Also called an Internet address.

IP multicast: Routing technique that allows IP traffic to be propagated from one source to a number of destinations or from many sources to many destinations. Rather than sending one packet to each destination, one packet is sent to a multicast group identified by a single IP destination group address. Packet replication occurs at routers in a classical router-based network. Multicast traffic uses class D IP addresses.

IR: insured rate. The long-term data throughput, in bits or cells per second, that an ATM network commits to support under normal network conditions. The insured rate is 100 percent allocated; the entire amount is deducted from the total trunk bandwidth along the path of the circuit.

IRTF: Internet Research Task Force. Community of network experts that consider Internet-related research topics. The IRTF is governed by the IRSG and is considered a subsidiary of the IAB.

ISDN: Integrated Services Digital Network. Communication protocol, offered by telephone companies, that permits telephone networks to carry data, voice, and other source traffic.

ITU-T: International Telecommunication Union Telecommunication Standardization Sector. International body that develops worldwide standards for telecommunications technologies. The ITU-T carries out the functions of the former CCITT.

LAN: local-area network. High-speed, low-error data network covering a relatively small geographic area (up to a few thousand meters). LANs connect workstations, peripherals, terminals, and other devices in a single building or other geographically limited area. LAN standards specify cabling and signaling at the physical and data link layers of the OSI model. Ethernet, FDDI, and Token Ring are widely used LAN technologies.

LANE: LAN emulation. Technology that allows an ATM network to function as a LAN backbone. The ATM network must provide some form of multicast and broadcast support, address mapping (MAC-to-ATM), SVC management, and a usable packet format. LANE also defines Ethernet and Token Ring ELANs.

LATA: local access and transport area. Geographic telephone dialing area serviced by a single local telephone company. Calls within LATAs are called "local calls," calls across LATA boundaries are "long distance" calls. There are over 100 LATAs in the United States.

LEC: LAN Emulation Client. Entity in an end system that performs data forwarding, address resolution, and other control functions for a single ES within a single ELAN. A LEC also provides a standard LAN service interface to any higher-layer entity that interfaces to the LEC. Each LEC is identified by a unique ATM address, and is associated with one or more MAC addresses reachable through that ATM address.

LEC: local exchange carrier. Local or regional telephone company that owns and operates a telephone network and the customer lines that connect to it.

LES: LAN Emulation Server. Entity that implements the control function for a particular ELAN. There is only one logical LES per ELAN, and it is identified by a unique ATM address.

LINF: Line Interface. Interface card used on the LightStream 100 ATM switch. The LINF receives cells sent over a line, checks them for errors, and forwards them toward their destination.

LightStream A-100: Cisco LightStream A-100 ATM switch is a fully-nonblocking ATM switch operating at up to 2.4 Gbps and supporting multiple ATM lines of 155-Mbps data speed as well as a variety of LAN and WAN interfaces. The LightStream A-100 switch can serve as part of an ATM workgroup or small campus backbone connecting a number of ATM routers, multilayer LAN switches, and high-performance servers and clients.

MB: maximum burst. Specifies the largest burst of data above the insured rate that will be allowed temporarily on an ATM PVC, but will not be dropped at the edge by the traffic policing function, even if it exceeds the maximum rate. This amount of traffic will be allowed only temporarily; on average, the traffic source needs to be within the maximum rate. Specified in bytes or cells.

MAN: metropolitan-area network. Network that spans a metropolitan area. Generally, a MAN spans a larger geographic area than a LAN, but a smaller geographic area than a WAN.

Mbps: megabits per second.

MCR: minimum cell rate. Parameter defined by the ATM Forum for ATM traffic management. MCR is defined only for ABR transmissions, and specifies the minimum value for the ACR.

Metasignaling: Process running at the ATM layer that manages signaling types and virtual circuits.

MTU: maximum transmission unit. Maximum packet size (in bytes) that a particular interface can handle. Classical IP uses a default MTU of 9,180 bytes. MTU for Ethernet LAN Emulation is 1,500 bytes. The maximum MTU for 10 Mbps Ethernet and for 100 Mbps Fast Ethernet is 1,500 bytes. The maximum MTU for 4 Mbps Token Ring is 4,472 bytes and for 16 Mbps Token Ring is 17,800 bytes. The maximum MTU for 100 Mbps FDDI is 4,500 bytes.

NHRP: Next Hop Resolution Protocol. Protocol used by routers to dynamically discover the MAC address of other routers and hosts connected to a NBMA network. These systems can then directly communicate without requiring traffic to use an intermediate hop, increasing performance in ATM, Frame Relay, SMDS, and X.25 environments.

NIC: network interface card. Board that provides network communication capabilities to and from a computer system. Also called an adapter.

N-ISDN: Narrowband ISDN. Communication standards developed by the ITU-T for baseband networks. Based on 64-Kbps B channels and 16- or 64-Kbps D channels.

NIST: National Institute of Standards and Technology. Formerly the National Bureau of Standards (NBS), this U.S. government organization supports and catalogs a variety of standards.

NMS: network management system. System responsible for managing at least part of a network. An NMS is generally a reasonably powerful and well-equipped computer such as an engineering workstation. NMSs communicate with software agents to help keep track of network statistics and resources.

NNI: Network-to-Network Interface. ATM Forum standard that defines the interface between two ATM switches that are both located in a private network or are both located in a public network. The interface between a public switch and private one is defined by the UNI standard. Also, the standard interface between two Frame Relay switches meeting the same criteria.

NSAP: network service access point. Network addresses, as specified by ISO. An NSAP is the point at which OSI Network Service is made available to a transport layer (Layer 4) entity.

NSF: National Science Foundation. U.S. government agency that funds a variety of scientific research in the United States.

OAM cell: Operation, Administration, and Maintenance cell. ATM Forum specification for cells used to monitor virtual circuits. OAM cells provide a virtual circuit-level loopback in which a router responds to the cells, demonstrating that the circuit is up, and the router is operational.

OC: Optical Carrier. Series of physical protocols (OC-1, OC-2, OC-3, and so on), defined for SONET optical signal transmissions. OC signal levels put STS frames onto multimode fiber-optic line at a variety of speeds. The base rate is 51.84 Mbps (OC-1); each signal level thereafter operates at a speed divisible by that number (thus, OC-3 runs at 155.52 Mbps). See Table 2.

OSI: Open System Interconnection. International standardization program created by ISO and ITU-T to develop standards for data networking that facilitate multivendor equipment interoperability.

OSI reference model: Open System Interconnection reference model. Network architectural model developed by ISO and ITU-T. The model consists of seven layers, each of which specifies particular network functions such as addressing, flow control, error control, encapsulation, and reliable message transfer. The highest layer (the application layer) is closest to the user; the lowest layer (the physical layer) is closest to the media technology. The lower two layers are implemented in hardware and software, while the upper five layers are implemented only in software. The OSI reference model is used universally as a method for teaching and understanding network functionality. Similar in some respects to SNA. In practice, network protocols do not always map clearly to the partitioned layers of the OSI Reference Model.

PCM: pulse code modulation. Transmission of analog information in digital form through sampling and encoding the samples with a fixed number of bits.

PCR: peak cell rate. Parameter defined by the ATM Forum for ATM traffic management. In CBR transmissions, PCR determines how often data samples are sent. In ABR transmissions, PCR determines the maximum value of the ACR.

PDU: protocol data unit. OSI term for packet. The term for packet varies across the OSI reference model, and the context of the discussion must be determined in order to determine what is being described. See Table 8 for a presentation of the layer descriptions.

Layer	Term for packet
Application	Message
Transport	PDU
Network	Datagram
Data Link	Frame, cell, packet
Physical	Frame

Table 8. OSI Layer Descriptions
for Packet.

PLCP: physical layer convergence procedure. Specification that maps ATM cells into physical media, such as T3 or E3, and defines certain management information.

PNNI: Private Network-Network Interface. ATM Forum specification that describes an ATM virtual circuit routing protocol, as well as a signaling protocol between ATM switches. Used to allow ATM switches within a private network to interconnect. Sometimes called Private Network Node Interface.

PSTN: Public Switched Telephone Network. General term referring to the variety of telephone networks and services in place worldwide. Sometimes called plain old telephone service (POTS).

PVC: permanent virtual circuit. Virtual circuit that is permanently established. PVCs save bandwidth associated with circuit establishment and tear down in situations where certain virtual circuits must exist all the time. Called a permanent virtual connection in ATM terminology.

PVP: permanent virtual path. Virtual path that consists of PVCs.

QoS: quality of service. Measure of performance for a transmission system that reflects its transmission quality and service availability. The ATM Forum has outlined five categories of performance (Classes 1 through 5) and recommends that ATM's quality of service should be comparable to that of standard digital connections.

QoS parameters: quality of service parameters. Parameters that control the amount of traffic the source router in an ATM network sends over an SVC. If any switch along the path cannot accommodate the requested QoS parameters, the request is rejected, and a rejection message is forwarded back to the originator of the request.

RARP: Reverse Address Resolution Protocol. Protocol in the TCP/IP stack that provides a method for finding IP addresses based on MAC addresses.

RBOC: Regional Bell Operating Company. Local or regional telephone company that owns and operates telephone lines and switches in one of seven U.S. regions. The RBOCs were created by the divestiture of AT&T. Also called Bell Operating Company (BOC). The RBOC structure is changing due to deregulation by the Telecommunication Act of 1996. Available at <http://thomas.loc.gov/cgi-bin/bdquery/z?d104:SN00652:>

RFC: Request For Comments. Document series used as the primary means for communicating information about the Internet. Some RFCs are designated by the IAB as Internet standards. Most RFCs document protocol specifications such as Telnet and FTP. A complete index of RFCs is available at <http://andrew2.andrew.cmu.edu/rfc/rfc-index.html>

rlogin: remote login. Terminal emulation program, similar to *telnet*, offered in most UNIX implementations. Passwords are not encrypted.

RPM: Reverse Path Multicasting. Multicasting technique in which a multicast datagram is forwarded out of all interfaces except the receiving interface, if the receiving interface is one used to forward unicast datagrams to the source of the multicast datagram.

RQ: rate queue. Value that is associated with one or more virtual circuits that defines the speed at which an individual virtual circuit will transmit data to the remote end. Each rate queue represents a portion of the overall bandwidth available on an ATM link. The combined bandwidth of all configured rate queues should not exceed the total bandwidth available.

SAP: service access point. Field defined by the IEEE 802.2 specification that is part of an address specification. Thus, the destination plus the DSAP define the recipient of a packet. The same applies to the SSAP.

SAR: segmentation and reassembly. One of the two sublayers of the AAL CPCS, responsible for dividing (at the source) and reassembling (at the destination) the PDUs passed from the CS. The SAR sublayer takes the PDUs processed by the CS and, after dividing them into 48-byte pieces of payload data, passes them to the ATM layer for further processing.

SCR: sustainable cell rate. Parameter defined by the ATM Forum for ATM traffic management. For VBR connections, SCR determines the long-term average cell rate that can be transmitted.

SDH: Synchronous Digital Hierarchy. European standard that defines a set of rate and format standards that are transmitted using optical signals over fiber. SDH is similar to SONET, with a basic SDH rate of 155.52 Mbps, designated as STM-1.

SDLC: Synchronous Data Link Control. SNA data link layer communications protocol. SDLC is a bit-oriented, full-duplex serial protocol that has spawned numerous similar protocols including HDLC and LAPB.

SDLC broadcast: Feature that allows a Cisco router that receives an all-stations broadcast on a virtual multidrop line to propagate the broadcast to each SDLC line that is a member of the virtual multidrop line.

SDU: service data unit. Unit of information from an upper-layer protocol that defines a service request to a lower-layer protocol.

SEAL: simple and efficient AAL. Scheme used by AAL5 in which the SAR sublayer segments CS PDUs without adding additional fields.

SMDS: Switched Multimegabit Data Service. High-speed, packet-switched, datagram-based WAN networking technology offered by the telephone companies. SMDS is based on IEEE standard 802.6.

SNI: Subscriber Network Interface. Interface for SMDS-based networks that connects CPE and an SMDS switch.

SNMP: Simple Network Management Protocol. Network management protocol used almost exclusively in TCP/IP networks. SNMP provides a means to monitor and control network devices, and to manage configurations, statistics collection, performance, and security.

ftp://ds.internic.net/rfc/rfc1157.txt — Simple Network Management Protocol (SNMPv1)
ftp://ds.internic.net/rfc/rfc1442.txt — Structure of Management Information for version 2 of the Simple Network Management Protocol (SNMPv2)

SONET: Synchronous Optical Network. High-speed (up to 2.5 Gbps) synchronous network specification developed by Bellcore and designed to run on optical fiber. STS-1 is the basic building block of SONET. Approved as an international standard in 1988.

SP: signaling packet. Generated by an ATM-connected device that wants to establish a connection with another such device. The signaling packet contains the ATM NSAP address of the desired ATM endpoint, as well as any QOS parameters required for the connection. If the endpoint can support the desired QOS, it responds with an accept message, and the connection is opened.

SSAP: source service access point. The SAP of the network node designated in the Source field of a packet.

SSCS: service specific convergence sublayer. One of the two sublayers of any AAL. SSCS, which is service dependent, offers assured data transmission. The SSCS can be null as well, in so-called Classical IP over ATM or LAN emulation implementations.

STM-1: Synchronous Transport Module level 1. One of a number of SDH formats that specifies the frame structure for the 155.52-Mbps lines used to carry ATM cells.

STS-1: Synchronous Transport Signal level 1. Basic building block signal of SONET, operating at 51.84 Mbps. Faster SONET rates are defined as STS-n, where n is a multiple of 51.84 Mbps.

STS-3c: Synchronous Transport Signal level 3, concatenated. SONET format that specifies the frame structure for the 155.52-Mbps lines used to carry ATM cells.

SVC: switched virtual circuit. Virtual circuit that is dynamically established on demand and is torn down when transmission is complete. SVCs are used in situations where data transmission is sporadic. Called a switched virtual connection in ATM terminology.

T1: Digital WAN carrier facility. T1 transmits DS-1-formatted data at 1.544 Mbps through the telephone-switching network, using AMI or B8ZS coding.

T3: Digital WAN carrier facility. T3 transmits DS-3-formatted data at 44.736 Mbps through the telephone switching network.

TAXI 4B/5B: Transparent Asynchronous Transmitter/Receiver Interface 4-byte/5-byte. Encoding scheme used for FDDI LANs as well as for ATM. Supports speeds of up to 100 Mbps over multimode fiber. TAXI is the chipset that generates 4B/5B encoding on multimode fiber.

T-carrier: TDM transmission method usually referring to a line or cable carrying a DS-1 signal.

TCP: Transmission Control Protocol. Connection-oriented transport layer protocol that provides reliable full-duplex data transmission. TCP is part of the TCP/IP protocol stack at the same level as best-effort UDP.

TCP/IP: Transmission Control Protocol/Internet Protocol. Common name for the suite of protocols developed by the U.S. DoD in the 1970s to support the construction of worldwide internetworks.

TDM: time-division multiplexing. Technique in which information from multiple channels can be allocated bandwidth on a single wire based on preassigned time slots. Bandwidth is allocated to each channel regardless of whether the station has data to transmit.

telnet: Standard terminal emulation protocol in the TCP/IP protocol stack. telnet is used for remote terminal connection, enabling users to log in to remote systems and use resources as if they were connected to a local system. Passwords are not encrypted.

TS: traffic shaping. Use of queues to limit surges that can congest a network. Data is buffered and then sent into the network in regulated amounts to ensure that the traffic will fit within the promised traffic envelope for the particular connection. Traffic shaping is used in ATM, Frame Relay, and other types of networks. Also known as metering, shaping, and smoothing.

trunk: Physical and logical connection between two ATM switches across which traffic in an ATM network travels. An ATM backbone is composed of a number of trunks.

TUD: trunk up-down. Protocol used in ATM networks that monitors trunks and detects when one goes down or comes up. ATM switches send regular test messages from each trunk port to test trunk line quality. If a trunk misses a given number of these messages, TUD declares the trunk down. When a trunk comes back up, TUD recognizes that the trunk is up, declares the trunk up, and returns it to service.

UBR: unspecified bit rate. QoS class defined by the ATM Forum for ATM networks. UBR allows any amount of data up to a specified maximum to be sent across the network, but there are no guarantees in terms of cell loss rate and delay.

UDP: User Datagram Protocol. Connectionless transport layer protocol in the TCP/IP protocol stack at the same level as reliable TCP. UDP is a simple protocol that exchanges datagrams without acknowledgments or guaranteed delivery, requiring that error processing and retransmission be handled by other protocols.

UNI: User-Network Interface. ATM Forum specification that defines an interoperability standard for the interface between ATM-based products (a router or an ATM switch) located in a private network and the ATM switches located within the public carrier networks. Also used to describe similar connections in Frame Relay networks.

UPC: usage parameter control. Process used to measure the actual traffic flow across a given connection and compare it to the total admissible traffic flow for that connection. Traffic outside of the agreed upon flow can be tagged (where the CLP bit is set to 1) and can be discarded en route if congestion develops. Traffic policing is used in ATM, Frame Relay, and other types of networks. Also known as admission control, permit processing, rate enforcement, and traffic policing.

URL: Uniform/Universal Resource Locator. Standardized addressing scheme for accessing hypertext documents and other services using a WWW browser. For example, the URL of this thesis is
<http://www.stl.nps.navy.mil/~iirg/atm/npslan/index.html>

UTP: unshielded twisted-pair. Four-pair wire medium used in a variety of networks. UTP does not require the fixed spacing between connections that is necessary with coaxial-type connections. There are five types of UTP cabling commonly used: Category 1 cabling, Category 2 cabling, Category 3 cabling, Category 4 cabling, and Category 5 cabling.

VBR: variable bit rate. QoS class defined by the ATM Forum for ATM networks. VBR is subdivided into a real-time (RT) class and non-real-time (NRT) class. VBR (RT) is used for connections in which there is a fixed timing relationship between samples. VBR (NRT) is used for connections in which there is no fixed timing relationship between samples, but that still need a guaranteed QoS.

VC: Logical circuit created to ensure reliable communication between two network devices. A virtual circuit is defined by a VPI/VCI pair, and can be either permanent (a PVC) or switched (an SVC). Virtual circuits are used in Frame Relay and X.25. In ATM, a virtual circuit is called a virtual channel.

VCC: virtual channel connection. Logical circuit, made up of VCLs, that carries data between two end points in an ATM network. Sometimes called a virtual circuit connection.

VCI: virtual channel identifier. 16-bit field in the header of an ATM cell. The VCI, together with the VPI, is used to identify the next destination of a cell as it passes through a series of ATM switches on its way to its destination. ATM switches use the VPI/VCI fields to identify the next network VCL that a cell needs to transit on its way to its final destination. The function of the VCI is similar to that of the DLCI in Frame Relay.

VCL: virtual channel link. Connection between two ATM devices. A VCC is made up of one or more VCLs.

VCN: virtual circuit number. 12-bit field in an X.25 PLP header that identifies an X.25 virtual circuit. Allows DCE to determine how to route a packet through the X.25 network.

VPC: virtual path connection. Grouping of VCCs that share one or more contiguous VPLs.

VPI: virtual path identifier. 8-bit field in the header of an ATM cell. The VPI, together with the VCI, is used to identify the next destination of a cell as it passes through a series of ATM switches on its way to its destination. ATM switches use the VPI/VCI fields to identify the next VCL that a cell needs to transit on its way to its final destination.

VPL: virtual path link. Within a virtual path, a group of unidirectional VCLs with the same end points. Grouping VCLs into VPLs reduces the number of connections to be managed, thereby decreasing network control overhead and cost. A VPC is made up of one or more VPLs.

WAN: wide-area network. Data communications network that serves users across a broad geographic area and often uses transmission devices provided by common carriers.

WWW or Web: World Wide Web. Large network of Internet servers providing hypertext and other services to terminals running client applications such as a Web browser. More generally, all information resources available via the Internet.

Web browser: Hypertext client application, such as Mosaic or Netscape, used to access hypertext documents and other services located on innumerable remote servers throughout the Web and the Internet. Text-based, graphical user interface (GUI) and 3D graphics browsers are available.

APPENDIX B. FORE CARD SPECIFICATIONS

A. INTRODUCTION

NPS purchased two Fore VMA-200 ATM VMEbus adapter cards that are compatible with Silicon Graphics workstations. The VMA-200 is a high performance adapter designed for use in a VMEbus/VME64-based system. The adapter provides ATM connectivity for VMEbus systems and is able to support evolving signaling and AAL standards. In addition, the VMA-200 provides transparent support for TCP/IP, SVCs through the SPANS signaling protocol (discussed later), PVCs, Q.2931 signaling, an ATM applications programmer interface (API), and an SNMP agent for network management. [Fore, 1995]

B. HARDWARE OVERVIEW

The VMA-200 is a single-slot VMEbus card supporting standard ATM cell processing including SAR. The card supports high quality image, full-motion video, CD-quality audio, and high speed data communications over a single ATM network connection. Figure 58 shows a view of the VMA-200 VMEbus adapter. [Fore, 1995]

C. SOFTWARE OVERVIEW

The VMA-200 uses Fore's proprietary signaling protocol, called simple protocol for ATM network signaling (SPANS), for establishing SVCs between other Fore system equipment. NPS purchased no Fore switches, so SPANS is not used. The VMA-200 card also provides an ATM Forum-compliant SNMP management information base




Figure 58. VMA-200 VMEbus Adapter, from [Fore, 1995].

(MIB) accessible by any SNMP-compatible network management system. Fore's application program interface (API) library is supplied with the VMA-200 card and offers applications access to ATM features such as guaranteed bandwidth reservation, per-connection selection of AAL5 or AAL3/4, and multicasting with dynamic addition and deletion of recipients. [Fore, 1995]

The VMA-200 card implements the ATM Forum's UNI 3.0 signaling specification (Q.2931 signaling) including support for interim local management interface (ILMI). The card also provides an implementation of classical IP using Q.2931 signaling. The classical IP implementation conforms to RFC 1577. [Fore, 1995]

D. VMA-200 ARCHITECTURE

The VMA-200 uses Fore's cell processing architecture, shown in Figure 59 .

Fore's strategy follows the conventional, high performance techniques — off-loading the communication processing from a central CPU. This architecture utilizes a dedicated embedded Intel i960 RISC processor along with special-purpose AAL5 and AAL 3/4 SAR hardware and scatter-gather DMA. The VAM is a VMEbus master device and supports 32 or 64 bit wide VMEbus data transfers and VMEbus DMA bursts of up to 64 bytes. [Fore, 1995]

E. HARDWARE REQUIREMENTS

The VMA-200 can be installed in an available VMEbus slot in any of the systems listed in Figure 59. Fore lists the device driver for the VMA-200 as multiprocessor safe. [Fore, 1995]

F. SOFTWARE REQUIREMENTS

The VMA-200 supports the IRIX operating system versions 5.3 and 6.0. Power Challenge and Power Onyx computers are supported with IRIX version 6.0. [Fore, 1995]

G. FIBER OPTIC CABLE SPECIFICATIONS

Table 9 lists Fore's recommendations for cable specifications. This is based on optimal adapter and switch performance.

H. TECHNICAL SPECIFICATIONS

The capabilities and physical parameter of the VMA-200 are detailed in Table 9.

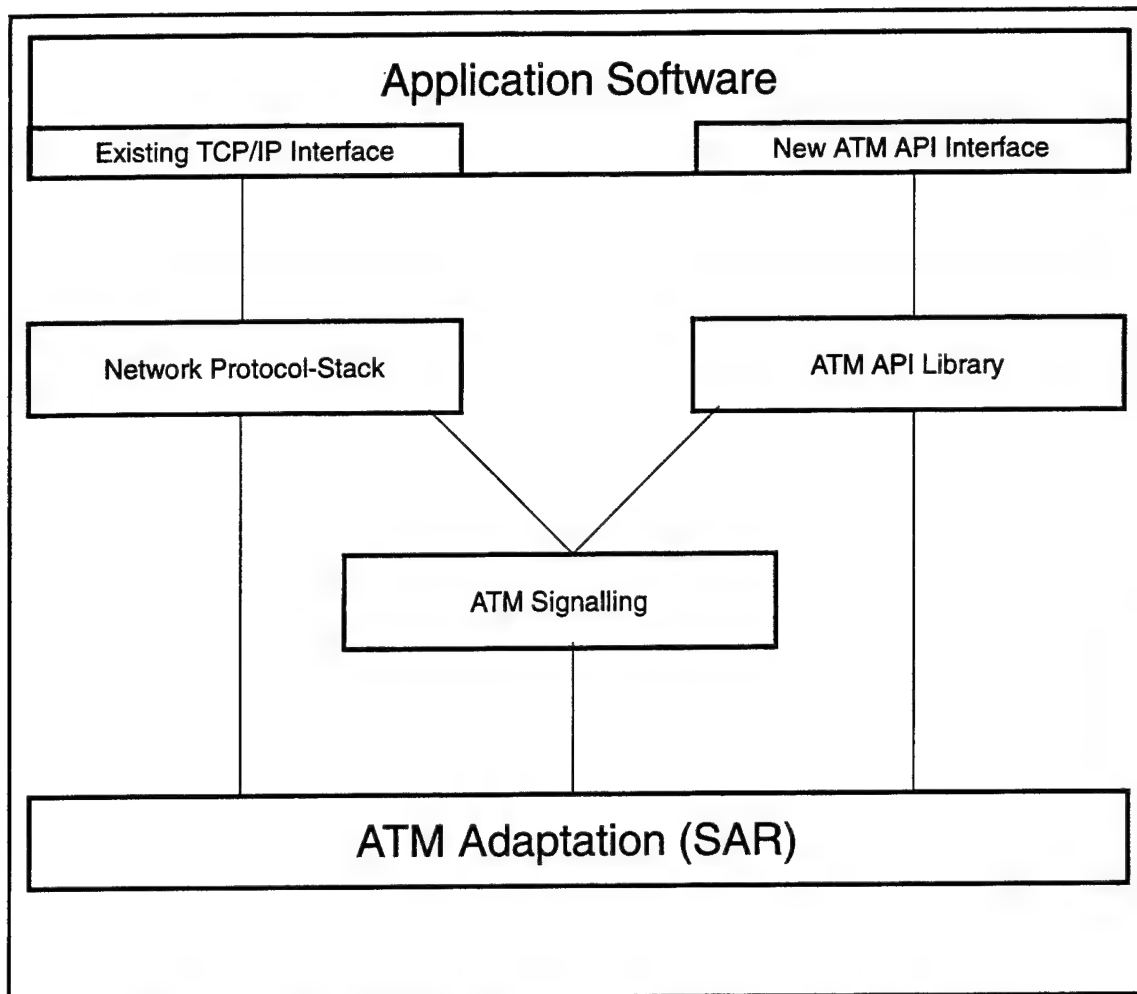


Figure 59. VMA-200 Cell Processing Architecture, from [Fore, 1995].

- Silicon Graphics, Inc., Challenge computer
- Silicon Graphics, Inc., Power Challenge computer
- Silicon Graphics, Inc., Crimson computer
- Silicon Graphics, Inc., Onyx computer
- Silicon Graphics, Inc., Power Onyx computer
- Any VMEbus/VME64 based system with customer supplied drivers

Figure 60. Systems That Can Use the VMA-200, from [Fore, 1995].

Description	Specification for Multi-Mode Products
Core Diameter	62.5 μm
Fiber Diameter	125 μm
Wavelength	1310 nm
Loss characteristic	~0.5 dB/km
Connector Style	SC or ST
Power Budget	11 dB (see note 1 below)
Approximate Distance	2 km
Transmit Power	-19 dBm (minimum)
Receive Power	-30 dBm (minimum)
Note 1 — If a 50 μm core fiber is used, the power budget is derated by 4 dB	

Table 9. Fiber Optic Cable Specifications for the VMA-200, from [Fore, 1995].

Hardware	
Architecture	On-board 25 MHZ i960 cell processor with VMEbus master burst transfer capability
AAL Support	Special purpose, on-board hardware for HEC and AAL5 and 3/4 calculations
UNI	100 and 140 Mbps TAXI (4B/5B encoding), OC-3/SONET
Form Factor	6U or 9U single-slot VMEbus
Compliance	ATM cell processing per ANSI T1S1.5/92-002R3, CCITT I.361, and ATM Forum v2.1 UNI specification
Cabling	Duplex 62.5/125 μ m multimode fiber (2000m max)
Connectors	ST
Software	
Transparent TCP/IP protocol interface	
Enhanced-performance ATM application programming interface (API) library	
SPANS SVC signaling protocol	
PVC	
Q.2931 signaling	
Support for ILMI	
Support for Classic IP interfaces	
Application-controlled multi/broadcasting with recipient add/delete capabilities	
SNMP MIB access to adapter status, ATM cell statistics, cell errors, and VCI/VPI	

Table 10. Technical Specifications of the Fore VMA-200, after [Fore, 1995].

APPENDIX C. CISCO A-100 SWITCH SPECIFICATIONS

A. INTRODUCTION

The discussion that follows of the Cisco LightStream A-100 switch is based on the technical specifications found in [Cisco, 1995].

The Cisco LightStream A-100 switch is a desktop ATM switch developed for workgroup LANs or campus backbone internetworks. The A-100 switch is an input and output buffer-type switch that supports 16 ATM lines at 155 Mbps, providing an aggregate throughput of 2.5 Gbps (16 lines x 155 Mbps/line). [Cisco, 1995]

The Cisco A-100 switch complies with the ATM Forum's *ATM User-Network Interface Specification, Version 3.0*, the International Telecommunication Union Telecommunication Standardization Sector's (ITU-T's) and the European Telecommunications Standards Institute's (ETSI's) specifications and recommendations. [Cisco, 1995]

The switch core is called an expandable ATM output-buffer modular switch (XATOMSW). The XATOMSW features large-capacity buffering to guarantee quality of service in handling multimedia traffic. [Cisco, 1995]

B. INTERFACE TYPES

The Cisco A-100 switch supports a wide range of LAN and WAN ATM interfaces. Different interface types can be mixed on an A-100 switch used for backbone, workgroup, or WAN access. Table 11 lists the interface types that are available. All

Physical Layer	Data Rate (Mbps)	Media	Connector
STS3c/STM1	155.52	Singlemode fiber	SC
STS3c/STM1	155.52	Multimode fiber	SC
STS3c/STM1	155.52	UTP-5	RJ-45
TAXI 4B/5B	100.00	Multimode fiber	MIC
DS3/T3	44.736	Coaxial cable	BNC
E3	34.00	Coaxial cable	BNC

Table 11. Cisco A-100 Switch Interface Types, from [Cisco, 1995].

interfaces conform to ATM standards, including those of the ATM Forum, ETSI, T1S1.5, and the ITU-T. [Cisco, 1995]

Because the Cisco A-100 switch is designed for workgroup and campus backbone deployment, it supports WAN interfaces such as DS3, E3, and single-mode fiber SONET. This capability allows seamless connectivity between an ATM campus backbone and ATM public and private WANs. In workgroups, the A-100 switch supports “power users” with direct ATM desktop interfaces by allowing copper (UTP-5) interfaces [Cisco, 1995]. Table 12 lists the ITU-T specifications that are supported, Table 13 lists the specifications of the switch, and Figure 61 lists the features of the switch.

Number	Topic
G.707	SDH speeds
G.708	SDH basic frame structure
G.709	SDH detailed frame structure
G.782	SDH multiplexer types and general characteristics
G.783	SDH multiplexer characteristics`
G.7XX	ATM cell mapping into SDH
G.803	Transmission network architecture
G.82X	Code error rules in physical layer
G.93B	B-ISDN subscriber line signal Layer 3
G.96X	B-ISDN digital section
I.113	B-ISDN terminology
I.121I	B-ISDN basic principles
I.150	B-ISDN ATM functionality
I.211	B-ISDN services
I.311	B-ISDN network side, signaling principles
I.321	B-ISDN protocol reference model and application
I.327	B-ISDN functional structure
I.35B	B-ISDN ATM layer cell transfer function
I.361	B-ISDN ATM layer specifications
I.362	B-ISDN AAL functions
I.363	B-ISDN AAL specifications
I.364	B-ISDN connectionless services
I.371	B-ISDN traffic control and congestion control
I.413	B-ISDN user network interface
I.414	ISDN and B-ISDN: concept of recommendation on Layer 1
I.432	B-ISDN user network interface Layer 1 specifications
I.580	B-ISDN and 64 kbps internetworking
I.610	B-ISDN OAM principles
Q.2931	B-ISDN: Signaling Layer 3
Q.2110	B-ISDN: AAL SSSCOP for signaling
Q.2130	B-ISDN: AAL SSCP for UNI signaling
Q.2140	B-ISDN: AAL SSCP for NNI signaling

Table 12. ITU-T Specifications Supported by the Cisco A-100, from [Cisco, 1995].

- Supports up to 16 155-Mbps ATM interfaces.
- Supports the addition of individual ATM interfaces of any physical layer type.
- Supports fully nonblocking, 2.5-Gbps input/output buffer-type switch fabric with a minimum of 1000 cells of virtual output buffering per port.
- Supports all AALs (AAL1 - AAL5) and traffic types.
- Supports two priority queues for cell delay: one for delay-sensitive traffic and one for delay-tolerant traffic. Cell loss priority is supported by configurable buffer threshold parameters.
- Provides fully integrated multicast capability without throughput degradation.
- Supports both PVCs and SVCs.
- Supports VC, VP, point-to-point, and point-to-multipoint connections and eliminates single points of failure through fully integrated support for ATM Forum UNI 3.0, based on Q.SAAL1 SSCOP.
- Allows construction of multiswitch networks using IISP and PNNI.
- Supports soft permanent virtual channel/permanent virtual path (PVC/PVP) connections.
- Supports SVC tunneling.

Figure 61. Cisco A-100 Switch Features, from [Cisco, 1995].

Device	Component	Specification
Switch	Switch architecture	Input and output buffer type
	Switch capacity	2.5 Gbps (155 Mbps x 16)
	Buffer	Input buffer: 2048 cells/2 lines
		Output buffer: 128 cells/2 lines
	Cell delay	20 μ sec to 5 msec
Control system	Control processor	Internal 32-bit RISC processor
	Number of concurrent connectable channels	4096 VP/VC channels/line
		All 12 bits of VPI
		Lower (12-x) bits of 16 bits of VCI
	NMS interface	SNMP
	Console terminal interface	EIA/TIA-232
	Ethernet interface	DB-15
Traffic control	Policing control	Peak cell rate can be set per connection
	Congestion control	Back pressure: output line \rightarrow switch \rightarrow input line
		Cell loss: two levels
	Priority control	Cell delay: two levels

Table 13. Cisco A-100 Switch Specifications, from [Cisco, 1995].

C. FUNCTIONAL OVERVIEW

The Cisco A-100 switch's three primary functions are listed in Figure 62.

1. Cell Switching Function

The Cisco A-100 switch's line interface (LINF) card receives a cell sent over a medium. A header translator performs the HEC and generates internal switch-specific overhead (SSO) information using VPI, VCI, and payload type (PT) in the cell header. If the cell belongs to a point-to-point connection, the LINF inserts new VPI and VCI values in the cell header. The SSO transfers to the XATOMSW along with the cell. The XATOMSW switches and sends the cell and SSO to the destination LINF according to the SSO information. When a specific line is congested, the system temporarily saves overflow cells in a 2048-cell input buffer and a 128-cell output buffer. Two lines share a buffer pool [Cisco, 1995]. A full discussion of switch construction and buffering considerations appears in [Varma, 1995].

The LINF inserts new VPI and VCI values in the cell header according to information in the SSO after receiving the cell and SSO if the cell belongs to a point-to-

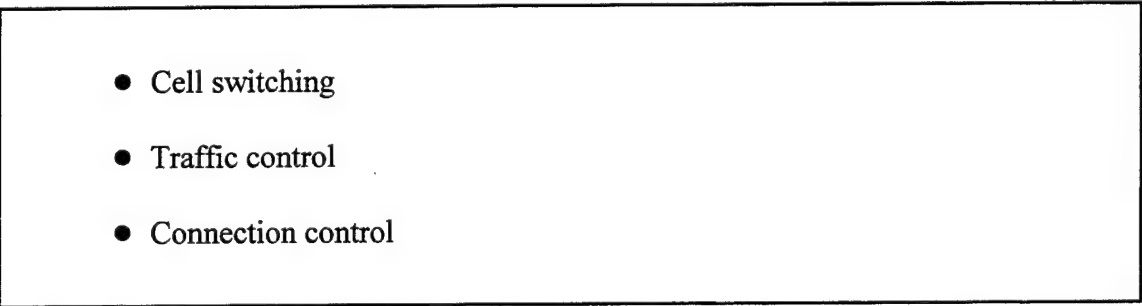
- 
- Cell switching
 - Traffic control
 - Connection control

Figure 62. Cisco A-100 Switch Primary Functions, from [Cisco, 1995].

multipoint connection. This series of events results in cell transmission to the destination line. [Cisco, 1995]

2. Traffic Control Functions

a. Policing Control

The Cisco A-100's line interface (LINF) card monitors each connection and discards any cell in excess of the preconfigured peak rate that is set in advance. This policing control, known as usage parameter control (UPC), prevents any connection from unjustly dominating the switch bandwidth. A peak transmission rate is set per connection. Once a peak rate is exceeded, the LINF discards excess cells if so configured. Specifically, the LINF counts the number of cells received per channel over a period of cell time (T) before the present moment. When the number of cells received over T exceeds a configured threshold (P), the system discards further cells. [Cisco, 1995]

The range of the cell time window (T) is 512 through 61,440. The cell time window (T) is divided into units of 512. Dividing the 61,440 window size into units of 512 yields 120 units. The measurement period can be specified in units of 1 through 120. Each unit represents a window time of 512 cells. The usage parameter value peak (UPVP) represents the number of cells within a measurement period. The peak value equals the UPVP times the number of units. The measurement period within the window is controlled by the *set tparam* command [Cisco, 1995]. Figure 63 shows the traffic measurement period. The constraint is given by Equation 3.




Figure 63. Traffic Measurement Period, from [Cisco, 1995].

$$\frac{\text{Transmission rate}}{\text{Link rate}} \leq \frac{\text{UPVP}}{512} \quad (3)$$

b. Congestion Control

In the Cisco A-100 ATM switch, a form of port congestion control takes place by applying a back pressure (BP) to restrict the number of input cells. The following format shows the direction that BP is applied:

Output line → Switch → Input buffer

Congestion control is discussed in detail later in this appendix.

c. Priority Control

Either high priority or low priority can be established for cell loss and cell delay on a per-connection basis. For cell delay, control is performed by the predetermined priority during the connection setup phase. For cell loss, control is performed by the value of the cell loss priority (CLP) bit in the cell header. The system

handles a cell for which the CLP bit in the header is set to 0 as high loss priority and during congestion will discard CLP bits set to 1 first. [Cisco, 1995]

A cell delay of 20 μ sec to 5 msec occurs in the switch, depending on the degree of congestion and degree of priority of the passing cells. Priority on the delay can be set according to the cell type as shown in Figure 64. [Cisco, 1995]

In the Cisco A-100 switch, two lines share one small input buffer pool (2048 cells). The maximum size allowed for guaranteed cells, best effort cells, and CLP threshold can be set in 128-cell increments for this input buffer. This setting applies to the entire A-100 switch. A different parameter cannot be set for each input buffer. [Cisco, 1995]

3. Connection Control

The A-100 switch performs resource allocation and path establishment when a PVC or SVC setup request occurs. The number of VPI and VCI bits necessary for an ATM connection is configured per line. A total of 12 bits is supported for VPI and VCI, up to 4096 channels per line. Because VPIs/VCI are distributed to PVCs and SVCs, when using SVCs the number of available PVCs diminished accordingly [Cisco, 1995].




Figure 64. Cell Type Priority, after [Cisco, 1995].

- The default number of allocated VPI/VCI bits is four for VPI and eight for VCI.
- The range of VCIs that can be assigned using eight bits is 1 to 255.
- 255 VCI values can be used for PVCs and SVCs.
- When using ITU-T, the four least significant VCI bits, 0 through 15, are reserved.
- When using ATM Forum, the five least significant VCI bits, 0 through 31, are reserved.

Figure 65. A-100 VPI/VCI Address Space, from [Cisco, 1995].

The A-100 supports 12 bits of VPI/VCI address space as discussed in Figure 65.

The A-100 switch can have both point-to-point and point-to-multipoint connections. Figure 66 indicates the number of possible PVC and SVC connection types per system. The number of point-to-point connections does not decrease if the maximum number of point-to-multipoint connections are established, and vice versa. [Cisco, 1995]

Entering commands at the console terminal can establish or disconnect

- Point-to-point PVC: 8,192
- Point-to-multipoint PVC: 1,024
- SVC: 1,024

Figure 66. A-100 PVCs and SVCs Possible, from [Cisco, 1995].

connections and can add endpoints on a point-to-multipoint connection [Cisco, 1995].

The A-100 switch controls the items listed in Figure 67 for each PVC connection. The details for establishing PVCs/SVCs are outlined in Appendix E.

D. EXPANDABLE ATM OUTPUT BUFFER MODULAR SWITCH (ATOMSW)

The ATOMSW functional block comprises eight multiplexer (MUX) /demultiplexer (DMUX) blocks and an ATOMSW. The ATOMSW is mounted on subboard, and the MUX/DMUX blocks are mounted on the motherboard [Cisco, 1995]. Figure 68 shows an ATOMSW block diagram. A full discussion of switch construction and the theory behind ATM switching and buffering appears in [Varma, 1995].

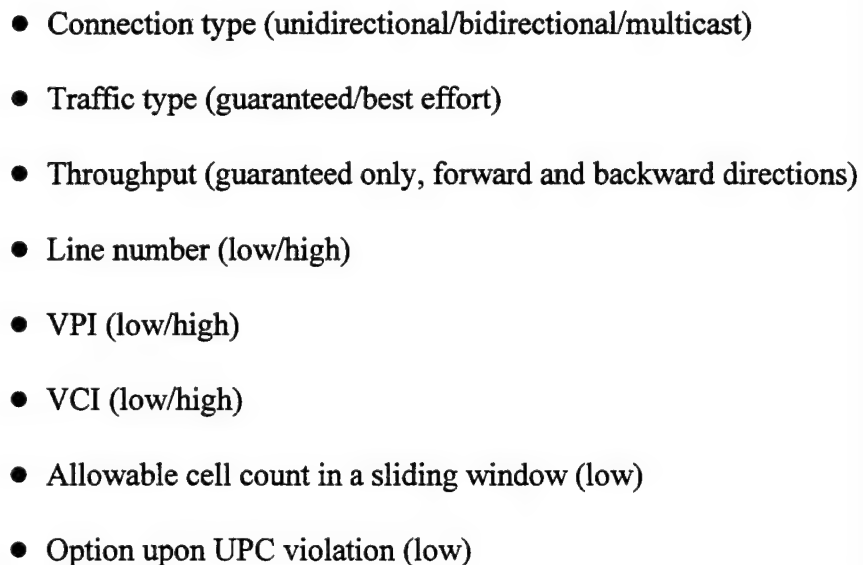
- 
- Connection type (unidirectional/bidirectional/multicast)
 - Traffic type (guaranteed/best effort)
 - Throughput (guaranteed only, forward and backward directions)
 - Line number (low/high)
 - VPI (low/high)
 - VCI (low/high)
 - Allowable cell count in a sliding window (low)
 - Option upon UPC violation (low)

Figure 67. PVC Elements Controlled by the A-100, from [Cisco, 1995].

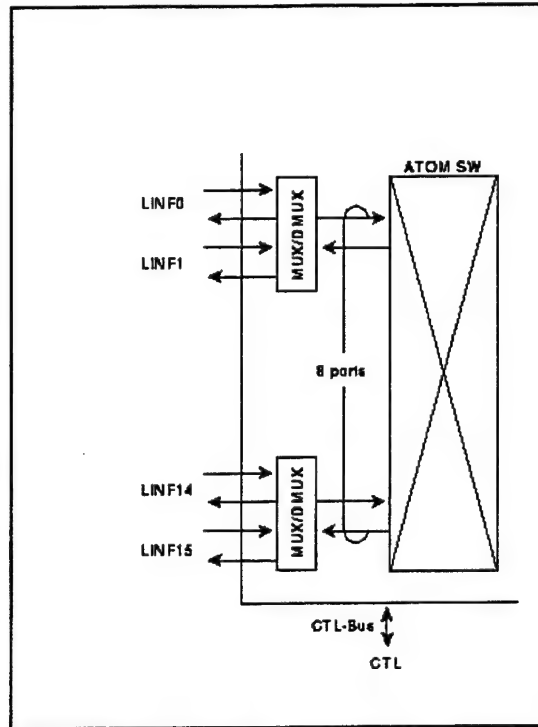
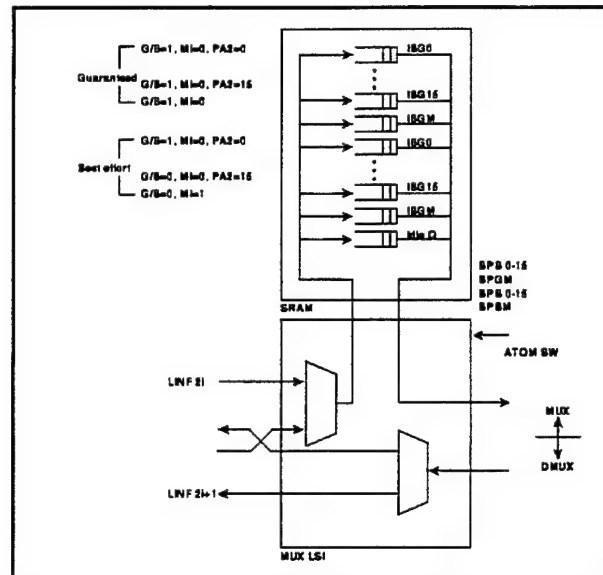


Figure 68. ATOMSW Block Diagram, from [Cisco, 1995].

1. MUX/DMUX

Each MUX/DMUX block consists of a MUX and a DMUX. The MUX statistically multiplexes two 155 Mbps cell streams into one time-division-multiplexed (TDM) 311 Mbps cell stream. The DMUX demultiplexes this 311 Mbps TDM stream into two 155 Mbps cell streams. The MUX has one input buffer (2048 cells) for every two lines. Figure 69 shows the MUX/DMUX card block diagram.

The MUX temporarily stores two 155 Mbps cell streams in a random-in, random-out (RIRO) buffer. The MUX then statistically multiplexes these into one 311 Mbps cell stream. The RIRO buffer comprises a total of 34 guaranteed and best effort first-in first-out (FIFO) cell queues: 32 FIFO queues for Line 0 — input buffer guaranteed 0 (IBG0)/



input buffer best effort 0 (IBB0) — through Line 15 (IBG15/IBB15); and two FIFO queues for multicast — input buffer guaranteed multicast (IBGM)/input buffer best effort multicast (IBBM). [Cisco, 1995]

The MUX multiplexes cells from two lines and temporarily stores multiplexed cells in the corresponding FIFO queue. The first cell of each FIFO queue is output when back pressure (BP) is not received from the destination line. From IBGM/IBBM, however, the first cell is transmitted when BP is not received from all lines. The cells in guaranteed FIFO queues are served before those in best effort queues [Cisco, 1995]. Table 14 provides the conditions in which these four FIFO type queues are served.

When the number of cells stored in each guaranteed or best effort FIFO queue exceeds a preset threshold, or when the number of cells stored in an idle queue goes below the preset threshold, the system discards the low-priority cells without saving them [Cisco, 1995].

2. ATOMSW

The ATOMSW is an output buffer-type cell switch with eight input ports and eight output ports. Figure 70 shows the ATOMSW block diagram. The buffer performs

Priority Level	IBGi	IBGM	IBBi	IBBM
1	Served when back pressure (BP) is not received and threshold is exceeded.			
2	Normal level		Served when BP is not received and threshold is exceeded.	
3			Normal level	
4 (cell output inhibited)	When BP is received		When BP is received	
Relevant Bps	BPGi	BPGi	BPGi	BPGi
		BPGM	BPGM	BPGM
			BPBi	BPBi
				BPBM

Table 14. Conditions in which FIFO Queues Are Served, from [Cisco, 1995].

9-bit (eight plus 1-bit for SSO) parallel processing by using bit-slice division and processing one bit with one ATOM buffer large scale integration (LSI). The ATOMSW resides on a subboard on the motherboard. Each output port is equipped with a 128-cell buffer. The CPU logically divides each output buffer into three FIFO cell queues. The first FIFO queue is a 48-cell queue for an even-numbered line, the next FIFO queue is a 48-cell queue for an odd-numbered line, and the last FIFO queue is a 32-cell queue for both lines. [Cisco, 1995]

a. ATOMSW Cell Switching Function

For a point-to-point connection, an output port is selected according to 8-bit routing information, called physical address 2 (PA2), in the switch-specific overhead (SSO). The four lower bits of the PA2 designate 16 lines. For a multicast connection, the

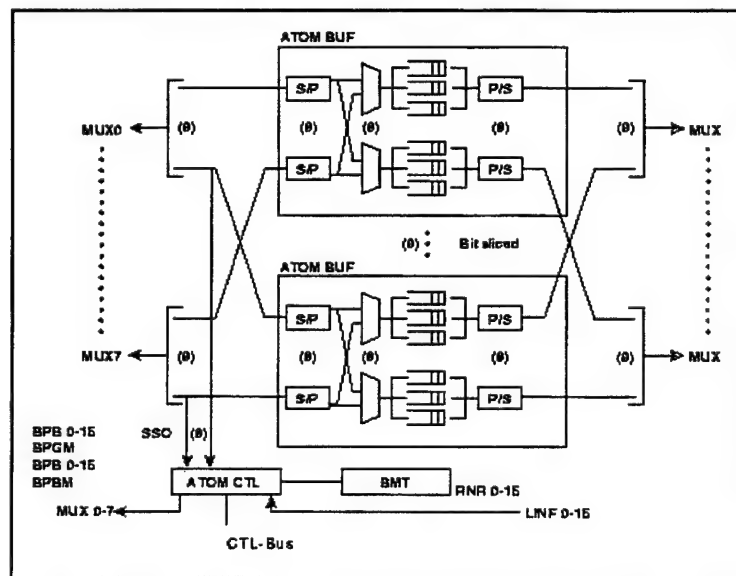


Figure 70. ATOMSW Block Diagram, from [Cisco, 1995].

multicast channel identifier in the SSO is referenced to acquire the bit map information on the output ports preset in the bit map table. This information is used to select output ports where the PA2 content is saved. [Cisco, 1995]

b. ATOMSW Congestion Control

Back pressure guaranteed (BPGi) is output when the number of cells stored in output line iFIFO queue exceeds a guaranteed threshold, and back pressure best effort (BPBi) is output when a best effort threshold is exceeded. Likewise, back pressure guaranteed multicast (BPGM) is output when the number of cells stored in multicast FIFO queue exceeds a multicast guaranteed threshold, and back pressure best effort multicast (BPBM) is output when a multicast best effort threshold is exceeded [Cisco, 1995]. Figure 71 lists these threshold values.

When a back pressure (BP) signal is output, up to two cells are input to each port, and cell input halts at every port. When BPi is received from output line I, cell output to line I halts. [Cisco, 1995]

- Guaranteed threshold: 32 (allowable range 0-48)
- Best effort threshold: 24 (allowable range 0-48)
- Multicast guaranteed threshold: 16 (allowable range 0-32)
- Multicast best effort threshold: 9 (allowable range 0-32)

Figure 71. ATOMSW Congestion Control Threshold Values, from [Cisco, 1995].

c. *ATOMSW Priority Control*

When the number of cells stored in the FIFO queue exceeds a threshold for priority control, the system discards the low-priority cells [Cisco, 1995]. Figure 72 lists these threshold values, and Figure 73 shows an ATOMSW output-buffer threshold values graph.

- Threshold: 40 (allowable range 0-48)
- Threshold for multicast: 24 (allowable range 0-32)

Figure 72. ATOMSW Threshold Values, from [Cisco, 1995].




Figure 73. ATOMSW Output-Buffer Threshold Values, from [Cisco, 1995].

APPENDIX D. FORE CARD CONFIGURATION

A. INTRODUCTION

Appendix D discusses the process for configuring the Fore card for ATM operation. Appendix E discusses the process for configuring the Cisco A-100 switch for ATM operation.

This Appendix discusses the NIC configurations that are installed for each phase of the NPS ATM LAN construction. The five step process is to install ATM in a stand-alone configuration, then peer-to-peer, through one switch, through two switches, and finally across campus and out to BayNet.

The Fore software for the VMA-200 ATM VMEbus adapter card is downloaded from two different sites. For ESA it is downloaded from

ftp.fore.com/priv/release/sunny/irix53_3.0.2e_1.2.tar.Z

and for VME from

ftp.fore.com/priv/release/sunny/irix53_3.0.2b_1.5.tar.Z

The software is installed in the */usr/etc/fore/etc* directories on Royal and Navy.

B. STAND ALONE

The stand alone configuration is shown in Figure 74. This configuration requires no special software configuration to be done since cells are not arriving over a medium.

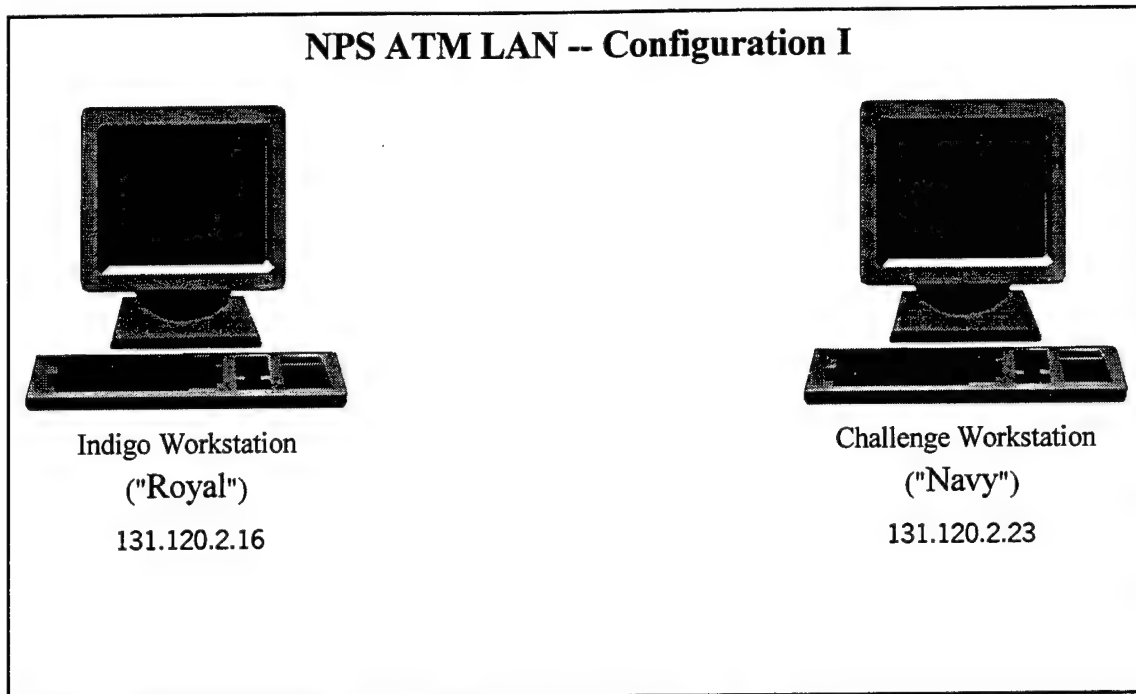


Figure 74. Stand Alone Configuration.

C. HOOKING UP TWO WORKSTATIONS PEER-TO-PEER

The peer-to-peer configuration is shown in Figure 75. The first thing that needs to be done is to configure the interface. When a Fore card talks to another Fore device, it uses a proprietary protocol over the fa0 interface. The commands are issued on Royal using the `ifconfig` (configure network interface parameters) command. The format for this command is `ifconfig interface IP address family subnet-mask broadcast addressess`.

```
ifconfig fa0 131.120.2.16 netmask 255.255.255.0 broadcast 131.120.2.255
rsh navy ifconfig fa0 131.120.2.23 netmask 255.255.255.0 broadcast
131.120.2.255
```

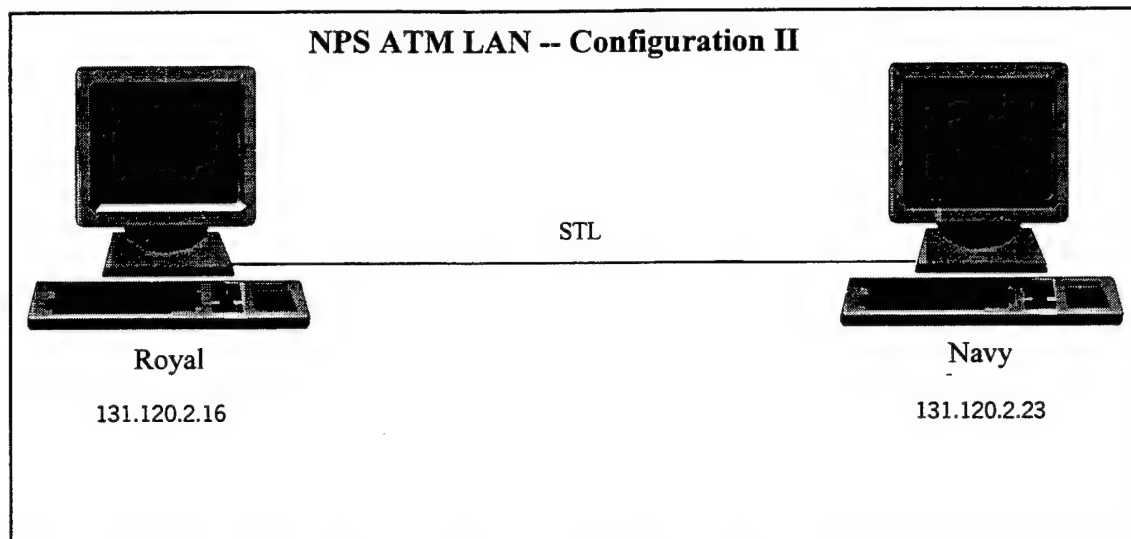


Figure 75. Peer-to-Peer Configuration.

The Fore help desk told us to turn off signaling and load balancing by issuing the -s and the -b switches, respectively, to the atmconfig command:

```
/usr/etc/fore/etc/atmconfig -s off fa0  
/usr/etc/fore/etc/atmconfig -b off  
rsh navy /usr/etc/fore/etc/atmconfig -s off fa0  
rsh navy /usr/etc/fore/etc/atmconfig -b off
```

Then setup the outgoing PVC on each machine. The command is atmarp -s (set permanent ARP entry for outgoing PVC) with the following format:

```
atmarp -s host device vpi vci aal
```

The commands are entered as follows:

```
/usr/etc/fore/etc/atmarp -s atm-navy fa0 0 100 5  
rsh navy /usr/etc/fore/etc/atmarp -s atm-royal fa0 0 100 5
```

Then setup the incoming PVC on each machine by following the same format using the -l switch:

```
/usr/etc/fore/etc/atmarp -l fa0 0 100 5
```

```
rsh navy /usr/etc/fore/etc/atmarp -l fa0 0 100 5
```

In order to remove the PVC later, the command lines that would be used are atmarp with the -d host switch (delete an ARP entry) and -x host device vpi vci switch (detach IP from an incoming PVC or SVC).

```
/usr/etc/fore/etc/atmarp -d atm-navy
```

```
/usr/etc/fore/etc/atmarp -x fa0 0 100
```

```
rsh navy /usr/etc/fore/etc/atmarp -d atm-royal
```

```
rsh navy /usr/etc/fore/etc/atmarp -x fa0 0 100
```

D. HOOKING UP A SINGLE ATM SWITCH TO TWO WORKSTATIONS

The single-switched configuration is shown in Figure 77. Since the same two workstations as before are being used, Navy and Royal, there is no need to

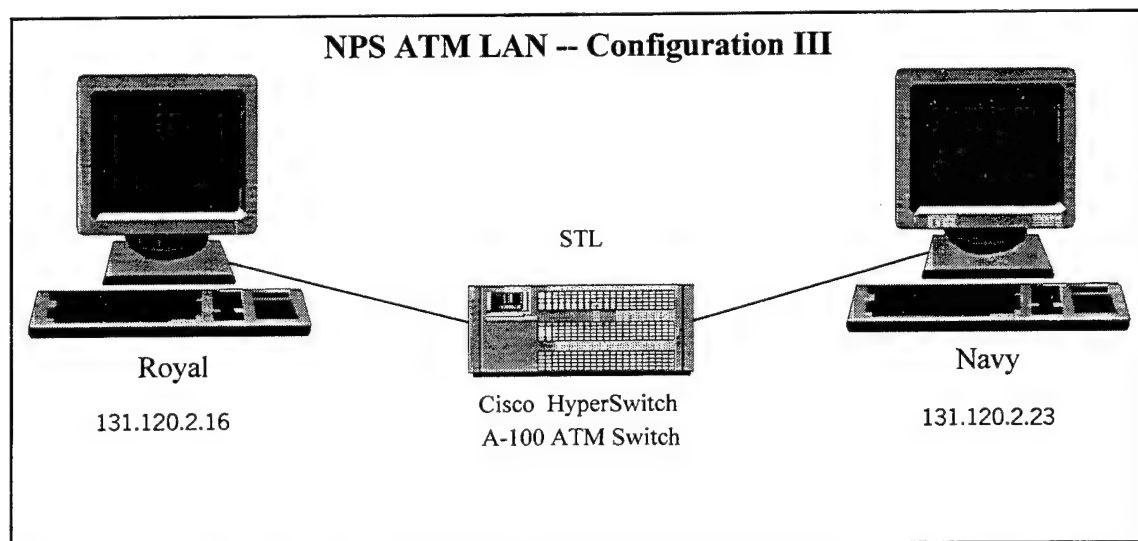


Figure 76. Single-Switched ATM LAN.

change any of the configurations for the Fore cards from the previous section. All changes at this point are in the Cisco A-100 switch. That is discussed in Appendix E.

E. HOOKING UP TWO ATM SWITCHES

The dual-switched configuration is shown in Figure 77. Since the same two workstations as before being used, Navy and Royal, there is no need to change any of the configurations for the Fore cards from the previous section. All changes at this point are in the Cisco A-100 switch. That is discussed in Appendix E.

F. CONNECTING TO BAYNET

The BayNet configuration is shown in Figure 78. Since the same two workstations as before are being used, Navy and Royal, there is no need to change any of the configurations for the Fore cards from the previous section. All changes at this point are in the Cisco A-100 switch. That is discussed in Appendix E.

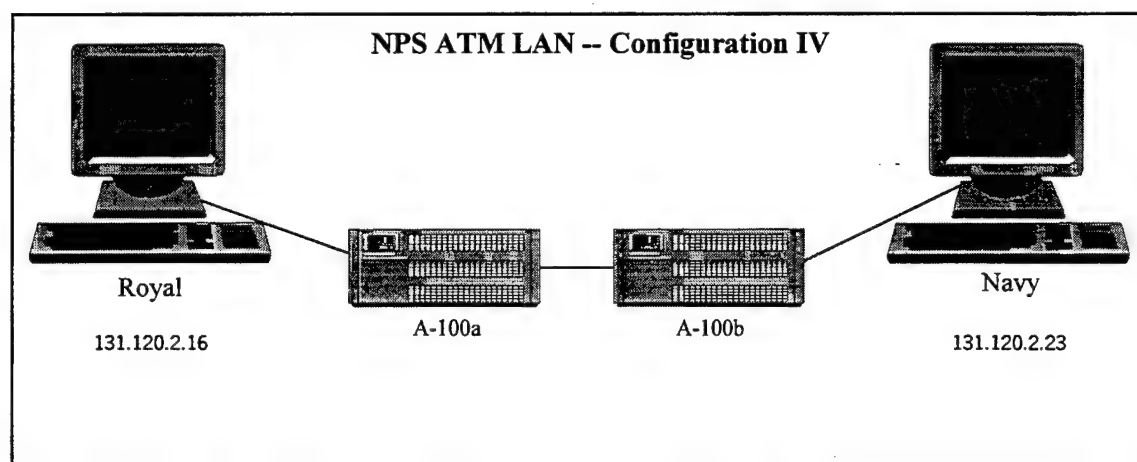


Figure 77. Dual-Switched ATM LAN.

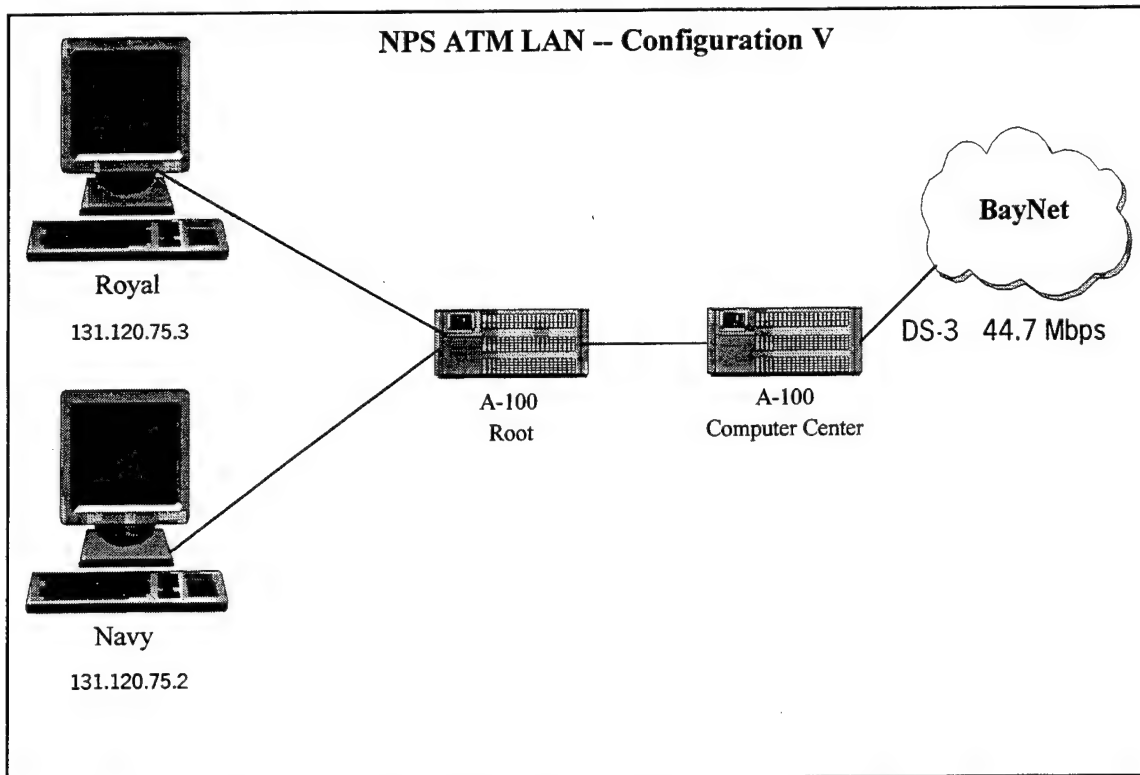


Figure 78. NPS ATM LAN Connection to BayNet Cloud.

APPENDIX E. CONFIGURING THE CISCO A-100 SWITCH

A. INTRODUCTION

Appendix D discusses the process for configuring the Fore card for ATM operation. Appendix E discusses the process for configuring the Cisco A-100 switch for ATM operation.

This Appendix discusses the configurations that are installed for each phase of the NPS ATM LAN construction. The five step process is to install ATM in a stand-alone configuration, then peer-to-peer, through one switch, through two switches, and finally across campus and out to BayNet.

B. STAND ALONE

The stand alone configuration is shown in Figure 79. The stand alone configuration requires no special software configuration since cells are not arriving by a switch.

C. HOOKING UP TWO WORKSTATIONS PEER-TO-PEER

The peer-to-peer configuration is shown in Figure 80. The peer-to-peer configuration required no special software configuration to be done since cells are not arriving by a switch.

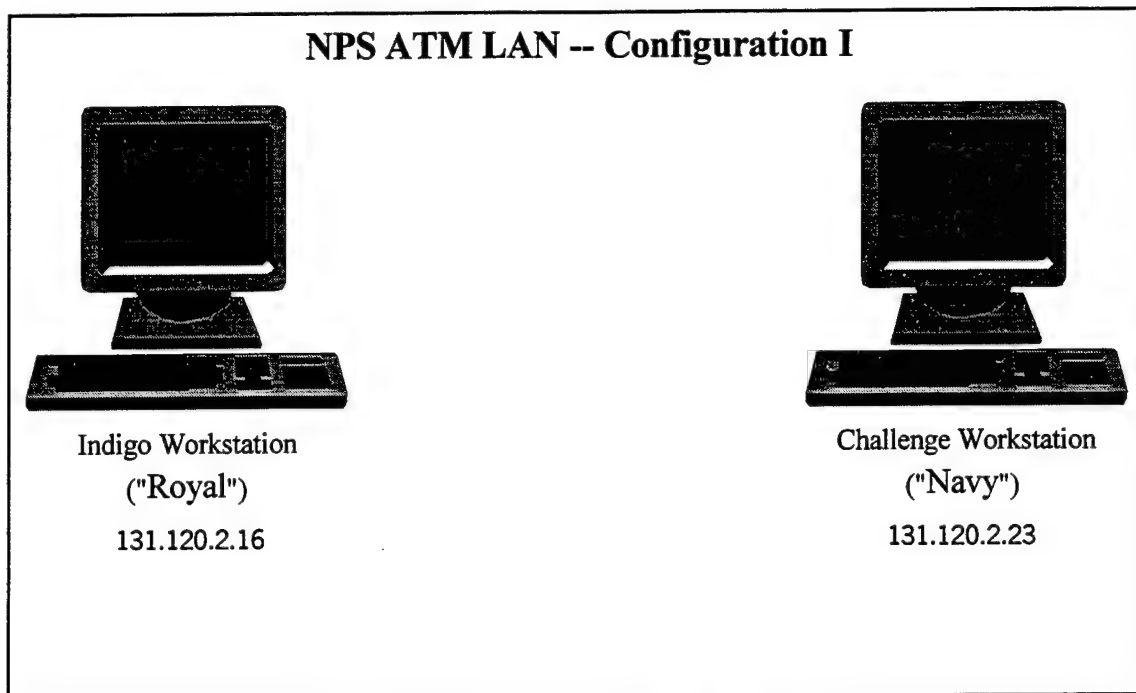


Figure 79. Stand Alone Configuration.

D. HOOKING UP A SINGLE ATM SWITCH TO TWO WORKSTATIONS

The single-switched configuration is shown in Figure 81. The fiber optic lines are plugged into ports 0 and 2 on the Cisco A-100 switch.

The switch console is connected from Royal to the ATM switch with a straight through cable from ttyd2 in order to talk with the switch and display the results. The command line switch is `cu` (call up another Unix terminal) to set the speed (`-s`) of 9600 bps and the device name (`-l`), `ttyd2`, to be used as the communication line.

```
cu -s9600 -l/dev/ttyd2
```

The line polarity is verified by issuing the following command to the A-100 switch:

```
show line
```

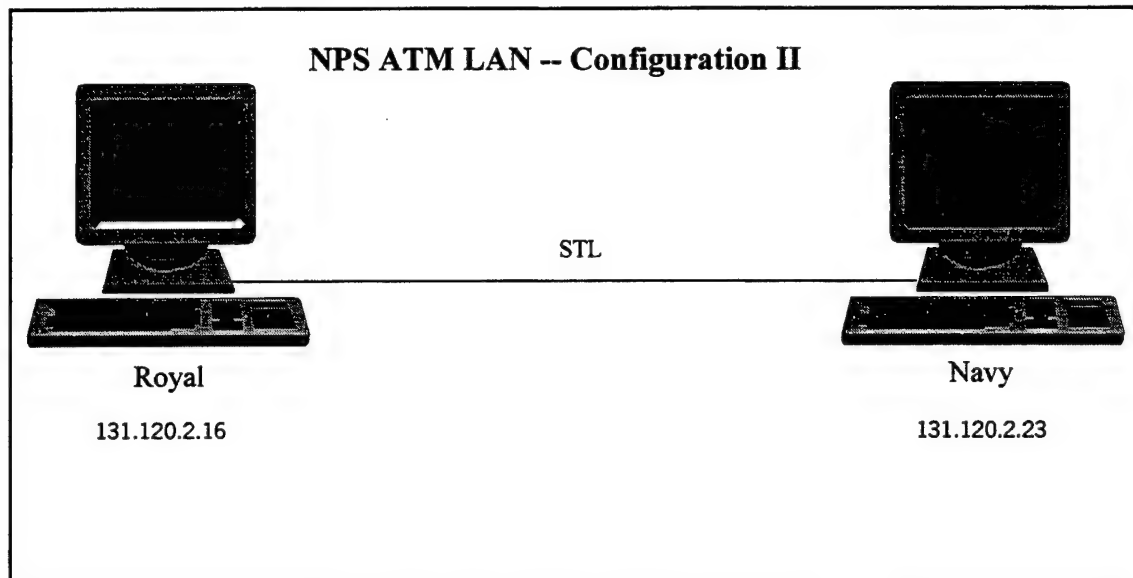


Figure 80. Peer-to-Peer Configuration.

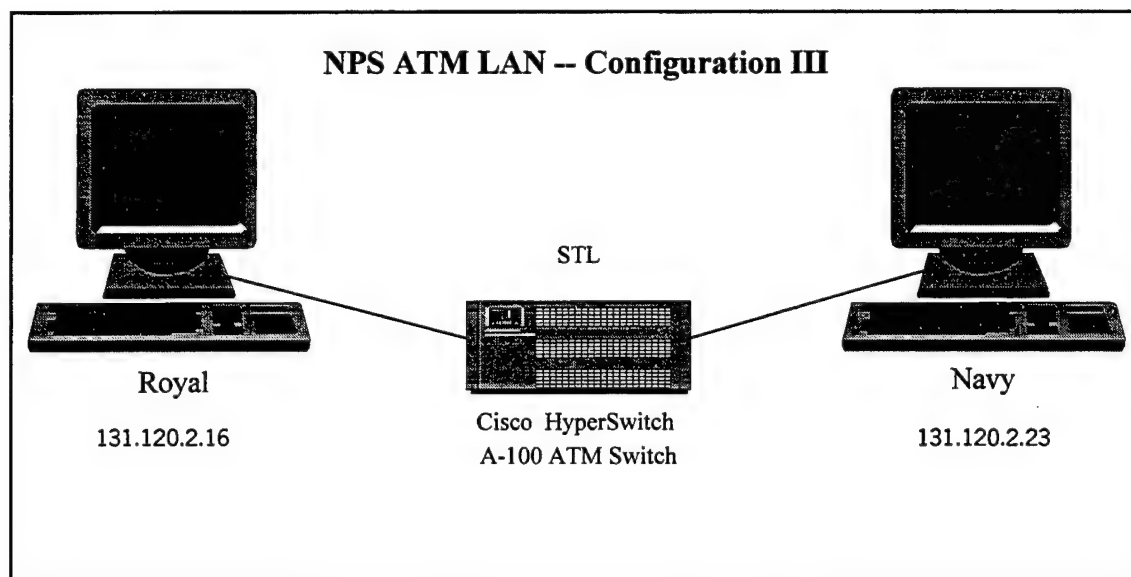


Figure 81. Single-Switched ATM LAN.

The output shows "Line 0 : (GOOD)". If it says "Loss of Signal" then the cables would need to be reversed on that port.

1. Configuring a PVC

The only step is to set up the switching for the PVC on the ATM switch. This is done using the `pvc establish` command which allows the user to specify the throughput (data rate) parameters for both forward and backward directions. `pvc establish` uses the following switches:

```
pvc establish p1 p2 p3 p4 p5 p6 p7 p8 p9 p10 p11 p12 p13 p14
```

Where,

- p1: Connection type (uni, bi, or multicast)
- p2: Traffic type (Guaranteed Service (GS), Best Effort (BE))
- p3: Low line number (0-15)
- p4: Low VPI (0-4095)
- p5: Low VCI (0-4095)
- p6: Low UPVP (0-512)
- p7: Low (forward) UPC enforcement option (through, discard)
- p8: Low (forward) line (allocated) transmission rate (Mbps)
- p9: High line number (0-15)
- p10: High VPI (0-4095)
- p11: High VCI (0-4095)
- p12: High UPVP (0-512)

p13: High (backward) UPC enforcement option (through, discard)

p14: High (backward) line (allocated) transmission rate (Mbps)

The following command is entered in order to setup a Bi-directional, Best Effort PVC with VCI 100 on line 2.

```
pvc establish 1 2 0 0 100 0 0 0 2 0 100 0 0 0
```

And the connection is completed.

2. Configuring an SVC

When a Fore card talks to another Fore device, it uses a proprietary protocol over the fa0 interface. The commands are issued on Royal using the `ifconfig` (configure network interface parameters) command. The format for this command is `ifconfig interface IP address family subnet-mask broadcast addresses`.

On both machines input:

```
ifconfig fa0 192.9.200.255
```

so that routes will be correct for qaa0.

On both machines input the following for the netmask for qaa0:

```
ifconfig qaa0 131.120.2.255 netmask 255.255.255.0
```

On both machines input the following to set up the qaa0 interface:

```
ifconfig qaa0 up
```

The NSAP address needs to be configured for qaa0 for both machines. This is done by using the `atmarp -n NSAPaddr dev` command. The NSAP address is aa followed by 38 zeros for Royal, and bb followed by 38 zeros . On Royal input:

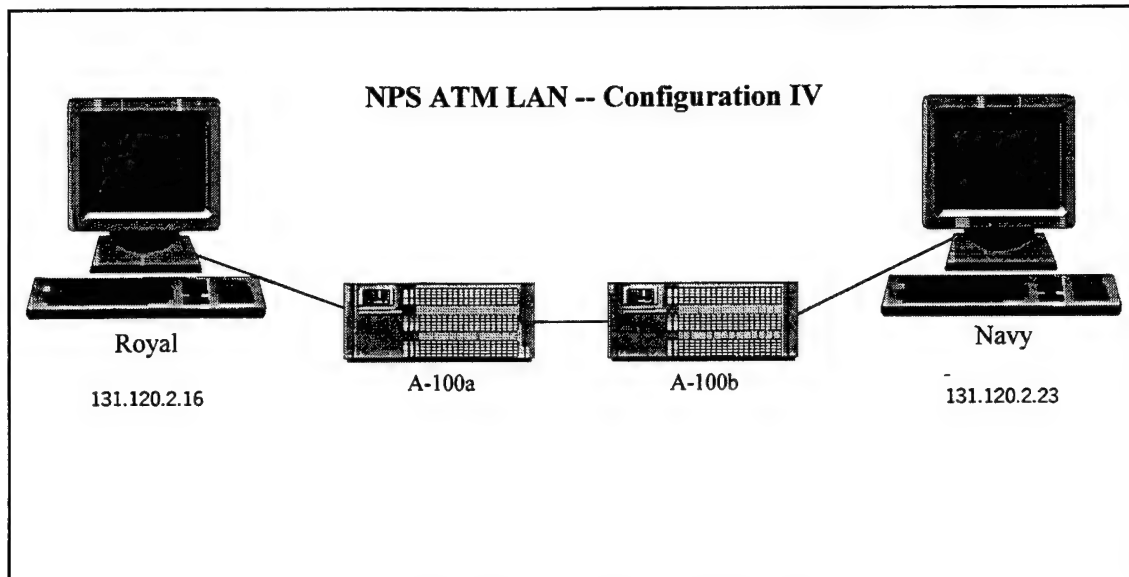


Figure 82. Dual-Switched ATM LAN.

On Royal:

```
ifconfig fa0 131.120.75.16 down
```

```
ifconfig qaa0 172.20.30.16 netmask 255.255.255.0 up
```

On Navy:

```
ifconfig fa0 131.120.75.23 down
```

```
ifconfig qaa0 172.20.30.23 netmask 255.255.255.0 up
```

Second, set NSAP addresses on the workstations:

On Royal

```
atmarp -n aa00000000000000000000000000000000 qaa0
```

On Navy

```
atmarp -n dd00000000000000000000000000000000 qaa0
```

The routes are checked on both hosts.

Qaa0 routes are added to both machines.

On Royal:

```
route add net 172.20.30 172.20.30.16 0
```

On Navy:

```
route add net 172.20.30 172.20.30.23 0
```

On CC's A-100 switch:

```
set local 172.20.30.1 255.255.255.0 0
```

```
route add dd0000 nsap 2
```

```
route add bb0000 nsap 0
```

```
route add aa0000 nsap 0
```

On STL's A-100 switch:

```
set local 172.20.30.2 255.255.255.0 0
```

```
route add aa0000 nsap 2
```

```
route add cc0000 nsap 0
```

```
route add dd0000 nsap 0
```

Add the entries to host ARP tables:

On Royal:

```
atmarp -o atm-royal aa00000000000000000000000000000000 qaa0
```

```
atmarp -o atm-navy dd00000000000000000000000000000000 qaa0
```

On Navy:

```
atmarp -o atm-royal aa00000000000000000000000000000000 qaa0
```

```
atmarp -o atm-navy dd00000000000000000000000000000000 qaa0
```

Upon trying to ping from Royal to Navy at this point nothing happens. The

traffic on ports 0 and 2 of the first switch are monitored while *pinging*. Data is making it through the first switch but the second switch shows that data is still not making it to the port 0. The Cisco trouble desk said to check the ATMSIG parameters of the A-100 switch:

```
show atmsig
```

The output is displayed in Figure 83.

On line (port) 0, one of the switches must be set to "User" and the other set to "Network" so that the switches can communicate with each other. The User/Network (U/N) switch in question for line 0 is in boldfaced type. To do this, issue the following commands to the computer center's A-100 switch:

```
set svcline 0 suspend
```

```
set atmsig 0 0 4 30 90 10 4 120 60 4 4 3.0 0
```

Line	U/ N	T303	T308	T309	T310	T313	T316	T317	T322	T398	T399	UNI Ver.	VPI
0	N	4	30	90	10	4	120	60	4	4	14	3.0	0
2	N	4	30	90	10	4	120	60	4	4	14	3.0	0
4	N	4	30	90	10	4	120	60	4	4	14	3.0	0
6	N	4	30	90	10	4	120	60	4	4	14	3.0	0
8	N	4	30	90	10	4	120	60	4	4	14	3.0	0
10	N	4	30	90	10	4	120	60	4	4	14	3.0	0

Figure 83. Output of the initial show atmsig command.

```
set svcline 0 resume
```

Then verify the correct configuration, the **show atmsig** command is issued again to the computer center's A-100 switch:

```
show atmsig
```

The results are displayed in Figure 84. U/N for line 0 has been set to U. *ping* now works with the SVC established on vpi.vci = 0.32.

F. CONNECTING TO BAYNET

The BayNet configuration is shown in Figure 85. UCSC provides the 172.20.70.0-net for this experiment. UCSC configures cyclone.cse.ucsc.edu as 172.20.70.2 and gives 172.20.70.1 for Royal to use. PacBell and Sprint set up a PVC between UCSC and NPS on vpi.vci=0.34. A classical IP PVC is used between the two

Line	U/ N	T303	T308	T309	T310	T313	T316	T317	T322	T398	T399	UNI Ver.	VPI
0	U	4	30	90	10	4	120	60	4	4	14	3.0	0
2	N	4	30	90	10	4	120	60	4	4	14	3.0	0
4	N	4	30	90	10	4	120	60	4	4	14	3.0	0
6	N	4	30	90	10	4	120	60	4	4	14	3.0	0
8	N	4	30	90	10	4	120	60	4	4	14	3.0	0
10	N	4	30	90	10	4	120	60	4	4	14	3.0	0

Figure 84. Output of the **show atmsig** command after resetting User/Network (U/N) status for line 0.

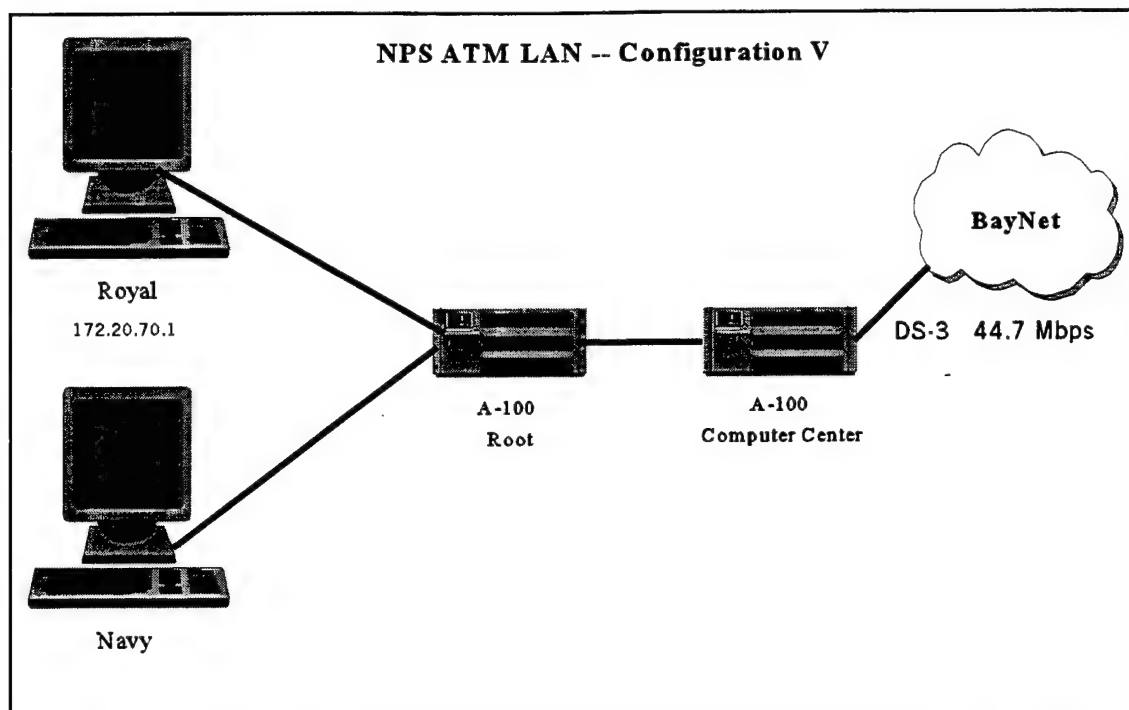


Figure 85. NPS ATM LAN Connection to BayNet Cloud.

machines, rate limited at the interfaces to 10 Mbps.

The first step is to set up the host table in the NIS:

```

172.20.70.1 atm-royal.stl.nps.navy.mil    atm-royal
172.20.70.2 cyclone.cse.ucsc.edu         nomad
  
```

On atm-royal the hostname, interface, and route are configured. Since only atm-royal will be used at first, atm-navy is unplugged for the test. The IP address of atm-royal is changed in its `/etc/host` file. Then the fa0 interface is brought down with another subnet number:

```
ifconfig qaa0 131.120.75.16 down
```

and the qaa0 interface is brought up with the UCSC subnet IP:


```
ifconfig qaa0 172.20.70.1 netmask 255.255.255.0 up
```

Routes are checked with `netstat -rn`, then a route is added for qaa0:

```
route add net 172.20.70 172.20.70.1 0
```

The switches are configured next. UCSC says that communication will initially be on vpi.vci=0.34, so a bidirectional PVC is established to/from port 0 on vpi.vci=0.34 to CC's A-100 switch, and to/from port 2 on vpi.vci=0.34 to atm-royal, by issuing the following commands to the STL A-100 switch:

```
PVC establish 1 2 0 2 100 0 0 0 2 0 100 0 0 0
save
```

The computer center then configures their switch similarly, but passes 0.34 traffic between their port 0 (DS3 out) and port 4 to STL. This is accomplished by issuing the following statements to CC's A-100 switch:

```
PVC establish 1 2 0 0 34 0 0 0 4 0 34 0 0 0
save
```

Next the interface cards are configured. No stray PVCs are verified by using the `atmarp -a` command. UCSC wants to limit traffic to 10 Mbps, so we set up the outgoing and incoming classical IP PVC on vpi.vci=0.34 limited to 10 Mbps:

```
/usr/etc/fore/etc/atmarp -c cyclone qaa0 0 34 0 10000
```

The preceding steps configure the entire PVC between Royal and Cyclone.

APPENDIX F. RETRIEVING ONLINE THESIS

This thesis is available at <http://www.stl.nps.navy.mil/~iirg/courtney/index.html>, and includes both hypertext markup language (HTML) as well as PostScript (PS) copies.

INITIAL DISTRIBUTION LIST

	No. Copies
1. Defense Technical Information Center 8725 John J. Kingman Road., Ste 0944 Ft. Belvoir, VA 22060-6218	2
2. Dudley Knox Library Naval Postgraduate School 411 Dyer Rd. Monterey, California 93943-5101	2
3. Don Brutzman, Code UW/Br Naval Postgraduate School Monterey, California 93943-5101	4
4. Rex Buddenberg, Code SM/Bu Naval Postgraduate School Monterey, California 93943-5101	2
5. Dale Courtney, Code 05A Naval Postgraduate School Monterey, California 93943-5101	2
6. Dr. Jim Eagle, Chair, Code UW Naval Postgraduate School Monterey, California 93943-5101	1
7. Dr. James C. Emery, Code 05 Naval Postgraduate School Monterey, California 93943-5101	2
8. Dr. Reuben Harris, Chair, Code SM/Ha Naval Postgraduate School Monterey, California 93943-5101	1
9. Dr. Ted Lewis, Chair, Code CS Naval Postgraduate School Monterey, California 93943-5101	1

10. Dr. G.M. Lundy, Code CS/Ln 1
Naval Postgraduate School
Monterey, California 93943-5101
11. Don McGregor, Code C3 1
Naval Postgraduate School
Monterey, California 93943-5101
12. Dr. Paul Moose, Code EC/Me 1
Naval Postgraduate School
Monterey, California 93943-5101
13. Dave Norman, Code 51 1
Director, W.R. Church Computer Center
Naval Postgraduate School
Monterey, California 93943-5101
14. Gary Porter, Code C3 1
Naval Postgraduate School
Monterey, California 93943-5101
15. Dr. Craig Rasmussen, Code MA/Ra 1
Naval Postgraduate School
Monterey, California 93943-5101
16. Dr. Maxine Reneker, Code 52 1
Director, Dudley Knox Library
Naval Postgraduate School
Monterey, California 93943-5101
17. Dr. Murali Tummala, Code EC/Tu 1
Naval Postgraduate School
Monterey, California 93943-5101
18. Terry Williams, Code C3 1
Naval Postgraduate School
Monterey, California 93943-5101
19. Commanding Officer, USS YORKTOWN (CG-48) 1
FPO AE 09599

20. Dr. Geoffrey Xie, Code CS/Xi 1
 Naval Postgraduate School
 Monterey, California 93943-5101

21. CAPT George Zolla, USN, Ret. 1
 Dudley Knox Library
 Naval Postgraduate School
 Monterey, California 93943-5101

22. Dr. Michael J. Zyda, Code CS/Zk 1
 Naval Postgraduate School
 Monterey, California 93943-5101

23. Arul K. Ananthanarayanan 1
 University of California Santa Cruz
 225 Applied Sciences Bldg
 Santa Cruz, California 95064

24. Dr. Larry Atkinson 1
 Center for Coastal Physical Oceanography
 Old Dominion University
 Norfolk, Virginia 23529

25. Dr. Marti Atkinson 1
 University of California Santa Cruz
 CIS/CE Board Office
 225 Applied Sciences
 Santa Cruz, California 95064

26. Roland Baker 1
 Santa Cruz County Office of Education
 Media and Technology Services
 809 Bay Avenue
 Capitola, California 95010

27. Lt Col Bill Barattino, Code JEEA 1
 Advanced Communications Technology Department
 Center for Systems Engineering
 Defense Information Systems Agency (DISA)
 10701 Parkridge Blvd
 Reston, Virginia 22091

28. CAPT Ross D. Barker 1
Senior Member, Smart Ship Program
Office of the Chief of Naval Operations
10701 Parkridge Blvd
Reston, Virginia 22091
29. Steve Batsell 1
Network Research Group Leader
MS 6367, Bldg 6012
Oak Ridge National Laboratory
Oak Ridge, Tennessee 37831
30. Dr. Carl R. Berman, Jr. 1
Office of the Provost
CSU Monterey Bay
100 Campus Center
Seaside, California 93955-8001
31. Bill Brutzman 1
3 South Kingman Road
South Orange, New Jersey 07079
32. Jeff Bryant 1
Monterey Bay Aquarium
886 Cannery Row
Monterey, California 93940
33. Erik Chaum 1
Code 2251, Building 1171
Naval Undersea Warfare Center Division Newport
1176 Howell Street
Newport, RI 02841-1708
34. Marke Clinger 1
FORE Systems, Inc.
174 Thorn Hill Road
Warrendale, PA 15086

35. Ray Cole 1
Naval Research Laboratory
Information Technology Division
Code 5523
b261/307a
4555 Overlook Ave.
Washington, DC 20375
36. Fred Cohn 1
Deputy to the City Manager
City of Monterey
City Hall
Monterey, California 93940
37. RADM O. Combs 1
SPAWAR-05
2451 Crystal Drive
Arlington, VA 22202
38. Coco Conn 1
2207 Willetta Avenue
Los Angeles, California 90068-2841
39. Dr. Tom Defanti 1
Electronic Visualization Laboratory
University of Illinois at Chicago, MC 154
851 S. Morgan St., Room 1120
Chicago, Illinois 60607-7053
40. Captain Dave Gamble 1
PMWW152
2451 Crystal Drive
Arlington, VA 22202
41. Dr. Nancy Giberson 1
Santa Cruz County Office of Education
809 Bay Avenue
Capitola, California 95010

42. Dr. Peter Giles, President and CEO 1
The San Jose Tech Museum of Innovation
145 West San Carlos Street
San Jose, California 95113
43. Mr. Paul Goss 1
SEA04I, Dir IRM
Naval Sea Systems Command
2531 Jefferson Davis Hwy.
Arlington, VA 22242
44. Bruce Gritton 1
Monterey Bay Aquarium Research Institute (MBARI)
P.O. Box 628
Moss Landing, CA 95039-0628
45. R. Peter Haddad, M/S 2L-3 1
Hewlett-Packard Laboratories
1501 Page Mill Road
Palo Alto, California 94304-1126
46. Dr. Kathryn Lee Hanson 1
Chief of Staff NII Activities
Silicon Graphics, Inc.
mail stop 001
2011 N. Shoreline Boulevard
Mountain View, California 94043-1389
47. Dr. Harold Hawkins 1
Office of Naval Research
Cognitive & Neural S&T Division
Program Officer
Attn: ONR 341
800 North Quincy Street, Ballston Tower One
Arlington, Virginia 22217-5660
48. Dr. G. Ross Heath 1
Executive Director
Monterey Bay Aquarium Research Institute (MBARI)
P.O. Box 628
Moss Landing, CA 95039-0628

49. Mike Herbst 1
Far West Laboratory
730 Harrison Street
San Francisco, California 94107-1241
50. Captain Roger Hull 1
PMW171
2451 Crystal Drive
Arlington, VA 22202
51. Birt Johnson 1
Pacific Bell
340 Pajaro Street, Room 131
Salinas, California 93901
52. Tom Karwin 1
U.C. Santa Cruz
P.O. Box 7600
Santa Cruz, California 95061-7600
53. Robert Kochanski 1
NRaD
Code 80
53560 Hull St.
San Diego, California 92152-5001
54. Dr. Marc Lipman 1
Mathematical, Computer, and Information Sciences Division
Office of Naval Research
Ballston Centre Tower One
800 North Quincy Street
Arlington, Virginia 22217-5660
55. Dr. Cathy Lascara 1
Center for Coastal Physical Oceanography
Old Dominion University
Norfolk, Virginia 23529

56. Dr. William J. "Bill" Lennon 1
Lawrence Livermore National Laboratory
7000 East Avenue (L-541)
Livermore, California 94550-9900
57. Syd Leung 1
Pacific Bell
2600 Camino Ramon, Room 35306
San Ramon, California 94105
58. Jeff Liening 1
FC4A/TNAC - Validation Support Branch
203 West Losey Street, Room 2000
Scott AFB, IL 62225-5238
59. Jack Linthicum 1
Technical Information Specialist
New Technology Development Division
Office of Engineering and Technology
Federal Communications Commission
1919 M Street N.W.,
Washington DC 20554
60. Brian Lloyd 1
Lloyd Internetworking
3461 Robin Lane
Cameron Park, California 95681
61. Cdr Mike Loescher 1
Spec Asst for IW
Deputy Assistant Secretary of the Navy (DASN)
Washington, D.C. 20350-1000

62. Dr. Melanie Loots 1
Associate Director, National Center for Supercomputing Applications (NCSA)
4119 Beckman Institute, MC 251
University of Illinois at Urbana-Champaign
605 East Springfield Ave.
Champaign, Illinois 61820
63. Dr. Michael R. Macedonia 1
Fraunhofer Center for Research in Computer Graphics, Inc.
Vice-president for developing global work environments
167 Angell Street
Providence, RI 02906
64. Rick Maciel 1
Pacific Bell ATM Project Manager
2600 Camino Ramon, Room 4S155
San Ramon, California 94583
65. Lora Lee Martin, Director 1
Director of Program and Policy Development
University of California, Santa Cruz
Santa Cruz, California 95064
66. Kam Matray 1
Monterey Bay Technology Education Center,
Monterey Peninsula Unified School District
PO Box 1031
Monterey, California 93942-1031
67. Chris May 1
California State University, Monterey Bay
100 Campus Center
Seaside, California 93955-8001
68. Dr. James H. May 1
California State University, Monterey Bay
100 Campus Center
Seaside, California 93955-8001

69. Michael McCann 1
Monterey Bay Aquarium Research Institute
P.O. Box 628
Moss Landing, California 95039-0628
70. Dr. Ray McClain 1
Moss Landing Marine Laboratories
P.O. Box 450
Moss Landing, California 95039
71. Mr. Jim McDonnell 1
SEA03KM
Naval Sea Systems Command
2531 Jefferson Davis Hwy.
Arlington, VA 22242
72. Mike Mellon 1
Monterey County Office of Education
Instructional Resources and Technology
PO Box 80851
Salinas, California 93912-0851
73. Dr. Pat Mantey 1
Chair, Computer Engineering
University of California, Santa Cruz
Santa Cruz, California 95064
74. Captain Mark Moranville 1
SEA03J, Dir Integrated IS Group
Naval Sea Systems Command
2531 Jefferson Davis Hwy.
Arlington, VA 22242
75. Michael Newman 1
Newman and Associates
24090 Summitwoods Drive
Los Gatos, California 95030

76. RADM(Sel) Kate Paige 1
 Commander NSWC
 10901 New Hampshire Ave
 Silver Spring, MD 20903-5640

77. Richard Pearlman 1
 Technical Director, Knowledge Network
 2150 Webster, Room 440
 Oakland, California 94611

78. Capt Jim Quiros 1
 FC4A/TNAC - Validation Support Branch
 203 West Losey Street, Room 2000
 Scott AFB, IL 62225-5238

79. Deborah Richards 1
 Industry Consultant, Pacific Bell
 2460 North Main Street
 Salinas, California 93906

80. Captain Charles Ristorcelli, PD16 1
 2451 Crystal Drive
 Arlington, VA 22202

81. Emily Routman 1
 Exhibit Developer, San Jose Tech Museum of Innovation
 145 West San Carlos Street
 San Jose, California 95113

82. Kathy Ruthowski 1
 Editor, NetTeach News
 13102 Weather Vane Way
 Herndon, Virginia 22071-2944

83. Dr. John Schill 1
 Program Manager
 ISO Planning and C3
 Defense Advanced Research Projects Agency
 3701 N. Fairfax Drive
 Arlington, Virginia 22203-1714

84. Dr. Arthur St. George 1
National Science Foundation
4201 Wilson Boulevard
Arlington, Virginia 22230
85. Dr. Fred Siff, Associate Vice Chancellor 1
Communications and Technology Services
University of California, Santa Cruz
1156 High Street
Santa Cruz, California 95064
86. G. (Joe) Simone 1
Executive Director,
FasTrak Advanced Digital Services
2600 Camino Ramon, Room 4S155
San Ramon, California 94583
87. Diane Siri 1
County Superintendent of Schools
Santa Cruz County Office of Education
809 Bay Avenue
Capitola, California 95010
88. Captain Ken Slaght, PMW176 1
2451 Crystal Drive
Arlington, VA 22202
89. Dr. Larry Smarr 1
Director, National Center for Supercomputing Applications (NCSA)
155 Computing Applications Building, MC 476
University of Illinois at Urbana-Champaign
605 East Springfield
Champaign, Illinois 61820
90. Bradley R. Smith, Computer Facilities Manager 1
University of California, Santa Cruz
145 Applied Sciences Bldg
Santa Cruz, California 95064

91. David Sonderegger 1
Pacific Bell ATM Project Manager
2600 Camino Ramon, Room 4S155
San Ramon, California 94583

92. LCDR Brian Steckler, USN, Ret. 1
25381 Carmel Knolls Drive
Carmel, California 93923

93. David Stihler 1
Monterey County Information Systems
1590 Moffett Street
Salinas, California 93905-3341

94. Chris Taylor 1
Director of Computing and Computer Resources (CCR)
California State University, Monterey Bay
100 Campus Center
Seaside, California 93955-8001

95. 2Lt Fred Taylor 1
FC4A/TNAC - Validation Support Branch
203 West Losey Street, Room 2000
Scott AFB, IL 62225-5238

96. Dr. Anujan Varma 1
Associate Professor of Computer Engineering
University of California, Santa Cruz
1156 High Street
Santa Cruz, California 95064

97. Jim Warner 1
Network Engineer, University of California Santa Cruz
C A T S/Network & Telco Services
11 Communications Bldg
Santa Cruz, California 95064

98. David Warren 1
Cabrillo College
6500 Soquel Drive
Aptos, California 95003

99. Marc W. Warshaw 1
124 White Oaks Lane
Carmel Valley, California 93924-9650
100. Steve Watkins 1
University of California Santa Cruz
Science Library
Santa Cruz, California 95064
101. Dr. Steve Webster 1
Director of Education
Monterey Bay Aquarium
886 Cannery Row
Monterey, California 93940
102. Dr. Glen H. Wheless 1
Center for Coastal Physical Oceanography
Old Dominion University
Norfolk, Virginia 23529